CHAPTER

7

NETWORKING

After reading this chapter and completing the exercises, you will be able to:

- Identify bus, ring, and star network topologies
- Discuss NetBEUI, IPX/SPX, and TCP/IP protocols
- Describe Token Ring and Ethernet media access methods
- Specify the purpose of bridges, switches, hubs, and routers
- List thinnet, shielded twisted-pair, unshielded twisted-pair, and fiber optic cable characteristics
- Make your own straight-through and crossover cables
- Describe network adapter teaming techniques
- Understand networking with a modem pool

Administrators encounter a myriad of network technologies and must be able to identify each one in order to take the best course of action in planning, extending, or troubleshooting the network. At the logistical center of every enterprise network is one or more servers; the goal of this chapter is to familiarize you with various ways to connect to that server.

As you analyze a network context, it is important to identify the network topology, protocol, and media access method in use, as each has advantages and disadvantages that affect network performance. Equally important is the choice of intermediate network equipment that directs data from one location on the network to another. Choosing the proper equipment in a given network context helps to ensure that servers have high availability.

NETWORK TOPOLOGIES

A network **topology** is the geometric configuration of devices, nodes, and cable links on a network. Topologies define how nodes connect to one another. A **node** is an active device connected to the network, such as a computer or a printer, or networking equipment such as a hub, switch, or router. (A host, defined in Chapter 1, is generally used interchangeably with node, but it is specific to devices using the TCP/IP protocol.) Nodes can be arranged in a bus, star, or ring configuration.

When discussing network topology, make sure you understand the difference between the physical topology and the logical topology. The physical topology is the layout of the actual connections between devices, while the logical topology is a representation of how data travels on the network. This distinction is observed throughout the chapter.

Before delving into the world of network topologies, let's look at bandwidth, which highlights several networking issues including topology, media access method, and physical media (cable).

Bandwidth

Recall that bandwidth is the transmission capacity of the network within a fixed amount of time. This is one of the fundamental factors that affect your choice of topology, media, and media access methods. Bandwidth has a direct correlation to the **data rate**, which is the actual quantity of data transferred within the limitations of the bandwidth. Bandwidth is usually expressed in bits per second (bps), kilobits per second (Kbps), kilobytes per second (KBps), megabits per second (Mbps), megabytes per second (MBps), gigabits per second (Gbps), gigabytes per second (GBps), terabits per second (Tbps), or terabytes per second (TBps). Note that the word "bits" is represented by a lowercase "b" and "bytes" (which is eight bits) by an uppercase "B."

For reference purposes, Table 7-1 shows most of the connection types with their associated data rates, bandwidths, and the time it takes to transfer 100 KB of data.

Bus

A **bus topology** consists of nodes linked together in a series where each node is connected to a common backbone cable or bus (see Figure 7-1). (A **backbone** is a larger, common avenue through which data transfers take place from smaller lines connected to it.) The signal is sent in both directions and has two endpoints (terminators) to prevent the signal from endlessly cycling through the cable. A major disadvantage of the bus topology is that it is more difficult to troubleshoot and locate a break in the cable or a faulty machine on a bus. A break anywhere in the cable will cause the entire segment to be inoperable until the break is repaired. Examples of bus topology include 10Base2 and 10Base5 Ethernet systems (addressed later in this chapter).

Table 7-1 Network Connection Types

Connection	Data Rate*	Bandwidth	Time per 100 KB
14.4 modem	1.8 KB	14.4 Kb	55 sec
28.8 modem	3.6 KB	28.8 Kb	27 sec
33.6 modem	4.2 KB	33.6 Kb	23 sec
56K modem	7 KB	56 Kb	14 sec
ISDN	7–16 KB	56–128 Kb	14–6 sec
Frame Relay	7–64 KB	56–512 Kb	14–1.5 sec
T-1	32–193 KB	256–1544 Kb	3.1–.5 sec
DSL	188 KB	1.5 Mb	.53 sec
Cable modem	188 KB	1.5 Mb	.53 sec
Fast Ethernet	1.25 MB	100 Mb	.08 sec
T-3	5.5 MB	44 Mb	.01 sec

* The actual data rate is slightly less than these figures because of overhead that occurs with framing the bits of data.

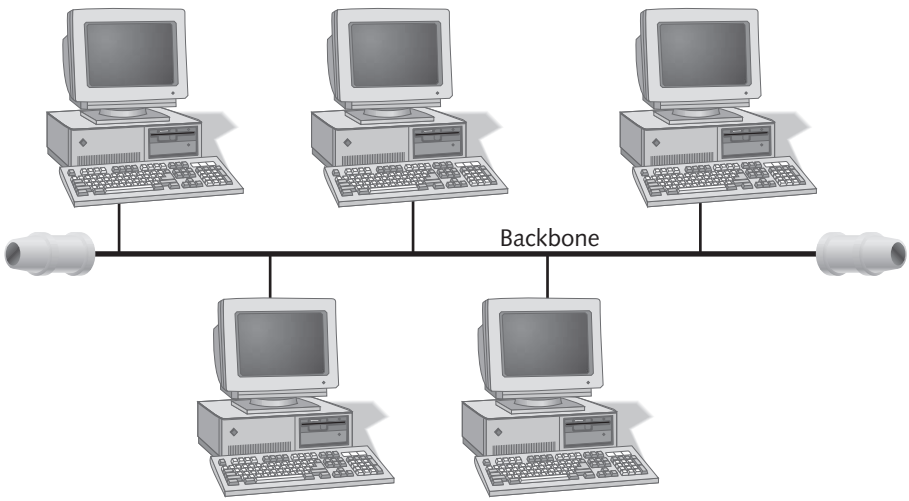

Figure 7-1 A backbone utilizes a bus topology to connect nodes

Ring

A **ring topology** network is a local area network (LAN) where all of the nodes are connected in a closed, single, logical communication loop (see Figure 7-2). Each device is connected directly to two other devices, one on either side. Information passes from station to station around the ring, each node reading the messages that are addressed to it and forwarding messages that are not. As with the bus network, each node must be able to identify its own address to successfully receive a message. A technique called

token passing manages line access so that two messages are not transmitted at the same time. A token is a frame of bits (the token may be "empty" or contain a message) that is passed from one station to the next. When a node needs to transmit data and receives an empty token, it holds on to the token and records its own address, the destination address, and the message into the token before passing it on to the next station. Stations only transmit messages when the token is empty.

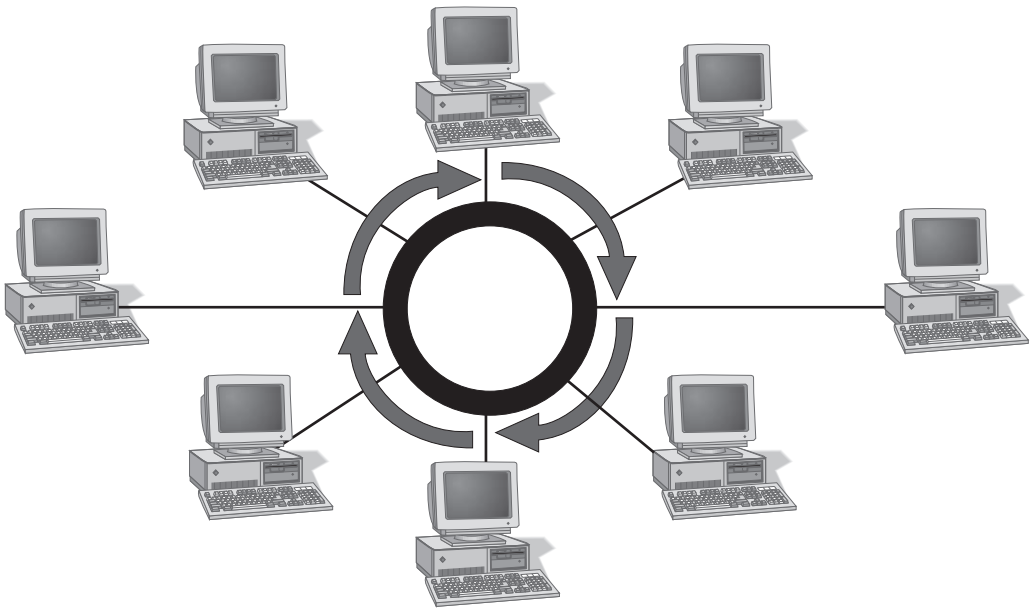

Figure 7-2 The ring topology

The destination station reads the message and consequently marks the token as having been read. The token passes from one node to the next until it completes a full circuit and reaches the originating station, where the message is discarded and the token is again marked as empty. Each node has a **transceiver**, which repeats the signal to move it around the ring.

Following are advantages of the ring topology:

- They can span greater distances than other network types—bus networks, for example.
- The level of signal deterioration is low because each station repeats the signal, and collisions are low because only the station that holds the token can transmit.
- It is very good for a small network of computers that entail high transmission speeds compared to 10BaseT Ethernet (addressed later in this chapter) or for larger networks where each station has a comparable workload.

Following are disadvantages of the ring topology:

- It can be tricky to trace a problem on the cable segment if the LAN is large.
- Each station's attached network interface must be continually active and the failure of a single station will halt a unidirectional ring network.
- It is complex to configure and requires relatively expensive hardware for each computer to interface with the network.
- Transmission delays tend to be long, even with moderate traffic levels.

Although a network design might utilize a logical ring topology as described here, the layout of the cables, nodes, and network equipment might physically be a star topology. (See the next section.)

Star

A **star topology** is a network configuration in which all of the nodes connect to a central network device such as a hub or switch (see Figure 7-3). All nodes receive the same signal, reducing effective bandwidth, and the central network device can become a bottleneck because all data must pass through it. Standard twisted-pair Ethernet networks using 10BaseT or 100BaseTX technology (addressed later in this chapter) commonly use the star topology.

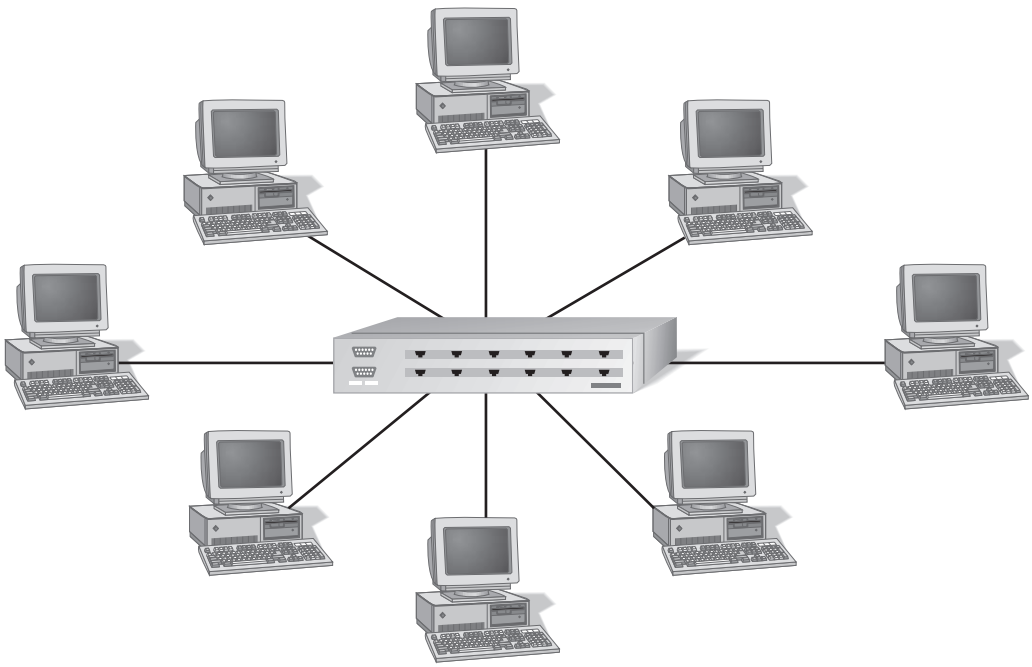

Figure 7-3 The star topology

Following are advantages of the star topology:

- A single failed node does not adversely affect the rest of the network.
- It is relatively straightforward to install and manage.
- Isolating and repairing bad segments is easier.
- It offers good capacity for network growth.

Following are disadvantages of the star topology:

- It requires a lot more cabling than bus or ring networks.
- The entire network becomes ineffectual if the central network device fails.

Hybrid Topologies

Network topologies are seldom of only one type. Except in the smallest environments where you could connect every user to a hub, topologies are usually a mixture. For example, a hybrid Ethernet network often uses a combination of the bus topology using either coaxial cable or fiber optic cable to connect multiple star-wired hubs. This creates a bus connection between the two hubs, while the hubs themselves are star-wired (see Figure 7-4). In the case of a Token Ring network, the network is a physical star and a logical ring, and various Token Ring networks can be connected together in a physical bus topology.

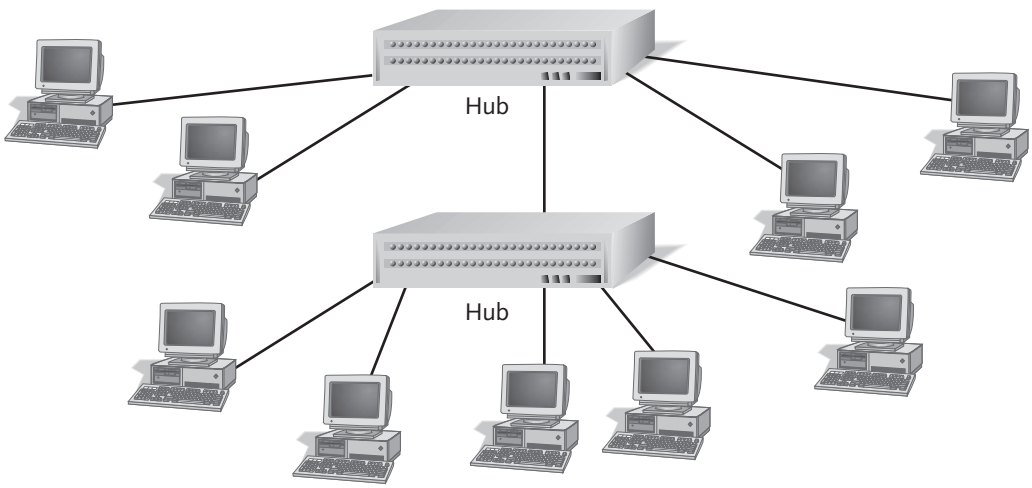

Figure 7-4 Ethernet networks commonly combine the bus and star topologies

There are several other hybrid network topologies, which are well covered in Chapter 5 of *Network+ Guide to Networks* by Tamara Dean (Course Technology, 2000).

PROTOCOLS

Recall from Chapter 1 that a protocol is a set of governing standards that determines how network devices communicate with one another. Also, a protocol defines how computers identify each other on a network, the form that the data should take in transition, and how this information is processed once it reaches its final destination. Protocols also define procedures for handling a lost or damaged packet—the electronic package that contains the network data. A protocol determines the type of error checking to implement, the data compression method (if any), how the sending device indicates that it is finished sending, and how the receiving device indicates that it is finished receiving.

There are several standardized protocols from which administrators can choose, each having its own particular advantages and disadvantages. Common protocols include NetBEUI, IPX/SPX, and TCP/IP. This chapter addresses TCP/IP more thoroughly than the other protocols because TCP/IP is more common and includes several utilities that you will certainly use to troubleshoot network connectivity and configuration.

Ultimately, all protocols are only a means to transport the network message to the node's physical address, known as the **MAC (Media Access Control) address**, which globally and uniquely identifies a network device. To find the MAC address of a network interface card (NIC), look at the MAC address printed on the NIC, or you can type IPCONFIG /all from a command prompt, which shows the MAC address as shown in the following example: Physical Address. . . . . : 00-03-47-12-39-FF. (See more about IPCONFIG later in this chapter.)

NetBEUI

NetBEUI is the **NetBIOS Enhanced User Interface**. Recall from Chapter 1 that NetBEUI is a fast protocol designed for small networks and requires no configuration. However, it is not a routable protocol and is not efficient in larger networks because it frequently rebroadcasts to locate other nodes on the network. It does not cache previously located nodes, and it does not use name resolution services such as DNS or WINS. (See more about DNS and WINS in Chapter 9.)

IPX/SPX

IPX/SPX (Internetwork Packet Exchange/Sequence Packet Exchange) is the default Novell protocol implementation for all versions of NetWare until 5.0, which can also use TCP/IP. IPX/SPX might require some configuration to identify the network on which the node exists, and like NetBEUI, it does not have name resolution services. However, IPX/SPX includes a caching mechanism so that it is not necessary to rebroadcast to locate recently accessed nodes. IPX/SPX was most popular when Novell NetWare networks required it, but current versions of NetWare can use TCP/IP instead, and many organizations are phasing out IPX/SPX.

TCP/IP

TCP/IP (Transmission Control Protocol/Internet Protocol) is actually a suite of protocols commonly in use on most networks and the Internet. However, administrators consider TCP/IP to be a single protocol. TCP/IP is more difficult to plan and configure, although it is also scaleable and routable, which is why it is the protocol of the Internet and most enterprise networks.

Covering all aspects of the TCP/IP protocol—including theory, configuration, and planning—is beyond the scope of this book. However, if you want to know more about TCP/IP, read any of dozens of comprehensive TCP/IP books.

The Internet (IP) Address

All protocols require a way to uniquely identify nodes. TCP/IP uses a unique IP address, similar to the way the U.S. Postal Service uses a combination of zip code, state, city, and street name to find its "nodes." An IP address appears as four sets of digits, each separated by a dot—215.161.122.231, for example. A host with this IP address must be unique on the LAN to avoid conflicts with other hosts, and it must be globally unique if the host IP address is exposed to the Internet. Each IP address also requires a **subnet mask**—another series of numbers which, when compared against the IP address, identifies the specific network to which the host belongs. As an example, Figure 7-5 shows a screen shot of an IP address and subnet mask for a Windows 2000 server.

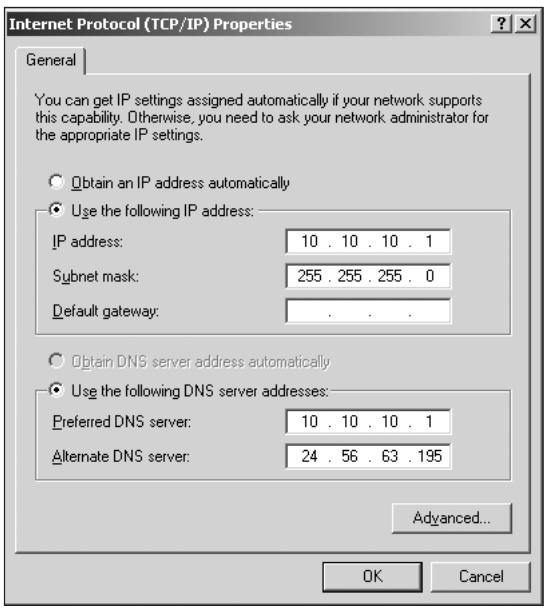

Figure 7-5 The TCP/IP configuration of a Windows 2000 server

Subnetting the Network

Administrators also use a subnet mask to divide a range of IP addresses into multiple smaller networks. The reason for doing this is twofold. First, you might not need all the IP addresses available on a single network. Instead, you can split the IP address range among several separate networks. Second, administrators subnet their networks to split up a collision domain, characteristic of Ethernet networks. (A **collision domain** refers to a network boundary in which multiple nodes could potentially attempt to access the network at the same time.) This chapter addresses Ethernet in more detail later, but for current purposes, understand that hosts on Ethernet networks have no arbiter to negotiate when network access is available (unlike the token of a ring network). As a result, network traffic grows exponentially as the number of nodes on a single network increases. Subnetting a larger network into smaller networks reduces the number of data collisions (and subsequent retransmissions) that take place when two hosts attempt to communicate at the same time. (A **collision** results when two devices or hosts transmit packets to the network at the same time.)

You can increase effective throughput to and from servers by installing network adapters with multiple ports, or multiple NICs in a single server. This makes the server **multihomed**. Similarly, you can use **port aggregation** software to combine multiple ports from the server into what is perceived as a single connection to the network but with bandwidth that is multiplied times the number of ports. Note that both multihoming and port aggregation only increase throughput of available bandwidth. The effectiveness of both methods diminishes with overutilized bandwidth.

Verifying TCP/IP Configuration and Connectivity

Connectivity or configuration problems with TCP/IP networks can involve lengthy and baffling troubleshooting. Fortunately, most network operating systems include a suite of TCP/IP configuration and troubleshooting tools to help diagnose problems.

Ping

Ping (packet internet groper) is an all-purpose utility for verifying that a remote host is accessible by sending small packets of data to which an accessible host responds. Ping tests connectivity at different stages between the host and destination to determine the point of failure at which a packet is dropped, and also tests basic networking connectivity. For example, an unplugged network cable would prevent Ping from reaching its destination, alerting you to a physical network problem (provided all TCP/IP configuration is correct).

In an IP network, Ping sends a single short data burst packet and then listens for a single packet in reply. Ping places a unique sequence number on each packet it transmits, and reports on the sequence numbers that come back to it. This is how Ping determines if packets have been dropped, duplicated, or reordered. Ping checksums (checks for

errors) in each packet it exchanges. Ping places a time stamp in each packet, which echoes back and computes the length of time for packet exchange. This is called the **round-trip time (RTT)**. Some routers silently discard undeliverable packets. Others mistake that a packet transmits successfully when it has not. Therefore, Ping may not always provide reasons why packets go unanswered. Ping does not perform analysis and cannot tell you why a packet was damaged, delayed, or duplicated. Ping also will not offer a play-by-play account of every host that handled the packet and everything that happened at every step of the way. Dropped packets are an unfortunate fact of networking life. There are common situations, typically involving crowded wide area networks (WANs), in which even modern TCP implementations cannot operate without dropping packets. Since TCP will retransmit missing data, there is no reason for alarm unless a large number of retransmissions noticeably affects network performance.

IPCONFIG

IPCONFIG is a Microsoft utility that displays a wide variety of IP configuration data for a Windows 98/ME/NT/2000 system including the IP address, subnet mask, default gateway, and other information (see Figure 7-6).

You can use the similar *netconfig* command for Linux/UNIX machines.

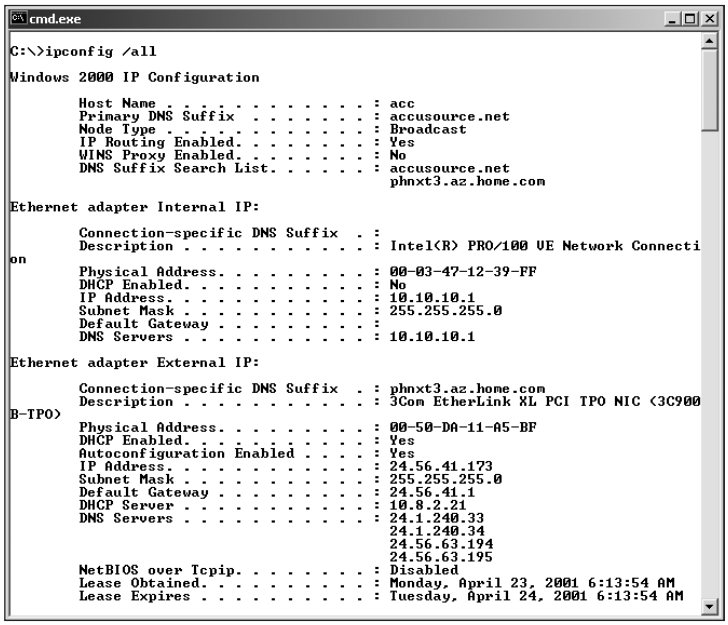

Figure 7-6 IPCONFIG /all displays complete IP configuration information

The most commonly used IPCONFIG switches are:

- /all—displays all available IP configuration information
- /release—releases IP configuration for DHCP clients
- /renew—renews the DHCP-assigned client IP address; useful when the DHCP configuration changes and you want to apply the changes to the DHCP client

ARP

ARP (Address Resolution Protocol) displays the resolution between the IP address and the physical (MAC) address on the NIC by building a table as IP addresses resolve to MAC addresses. You can also modify the ARP cache and table entries. For example, you can use the –s switch to add a static host-to-MAC address entry to the ARP table. The advantage of this would be to improve IP-address-to-MAC-address resolution time. To view the ARP table, type ARP –a, and to add an ARP entry, type ARP –s <IPAddress> <MAC address> (see Figure 7-7).

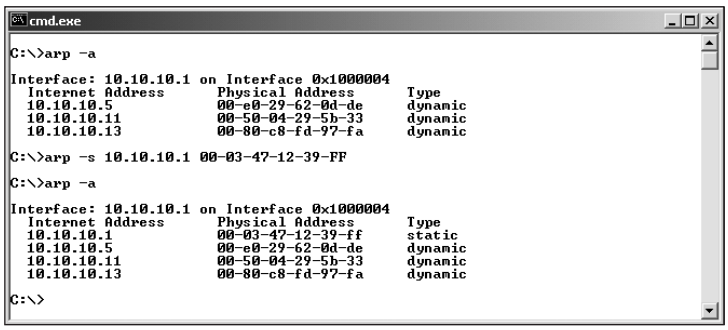

Figure 7-7 Use ARP for IP-address-to-MAC-address issues

TRACERT

TRACERT is the trace routing utility that works like Ping but shows the actual router hops taken to reach the remote host. This is handy if you want to find at which point a packet is being dropped or where a bottleneck may exist on the network (see Figure 7-8).

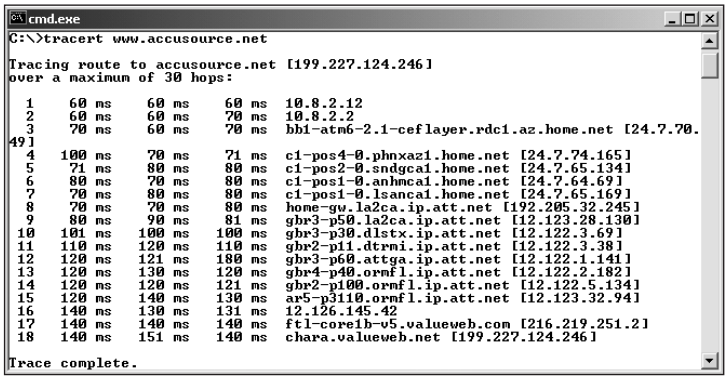

Figure 7-8 TRACERT identifies each router hop

NETSTAT

NETSTAT shows TCP/IP protocol statistics using any of several options. One of the most useful options is –r, which shows the routing table (see Figure 7-9). This is useful in verifying the efficiency of the routing tables.

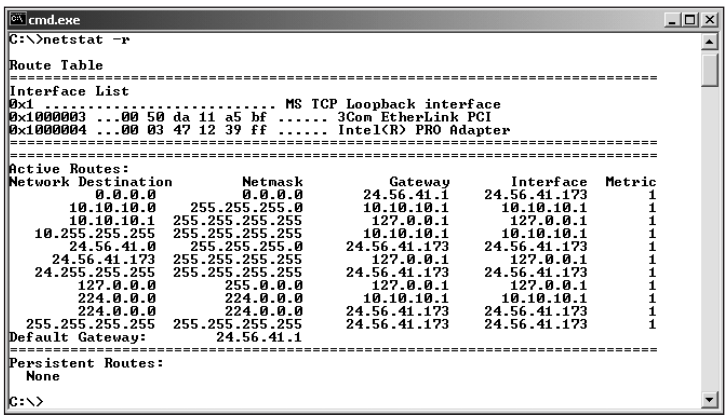

Figure 7-9 NETSTAT -r shows the routing table

NETWORK MEDIA ACCESS METHODS

Network communication requires a **media access method**, a way to place the data packets transmitted from the NOS to the physical network device (such as a NIC) and then to the wire. Several media access methods fulfill this role, each having its own characteristics.

Token Ring

Token Ring is a type of network where all of the computers are arranged in a circle. A special bit pattern, called a token, moves around the circle. To send a message, a station grabs the token, affixes a message to it, and then allows it to continue around the ring network. As a network protocol like Ethernet, Token Ring refers to the PC network protocol developed by IBM that has been standardized with the IEEE 802.5 standard. Token Ring is different from Ethernet in that all messages are transferred in one direction along the ring at all times. Numerous PC vendors have been proponents of Token Ring networks at different times; therefore, these types of networks can be found in many organizations.

Token Ring networks run at 4 or 16 Mbps. If a 16 Mbps adapter exists on a network where 4 Mbps adapters exist, you must configure the 16 Mbps adapter to run at the slower 4 Mbps speed. If you mix the two speeds on the same network, serious communication problems occur.

As the name implies, the Token Ring media access method utilizes a logical ring topology. However, the physical layout is usually a star topology with each host connecting to a **multistation access unit (MAU)**, which looks much like a hub except that it includes an RI (ring in) and RO (ring out) port. Tokens still pass from one host to the next in a logical ring (see Figure 7-10).

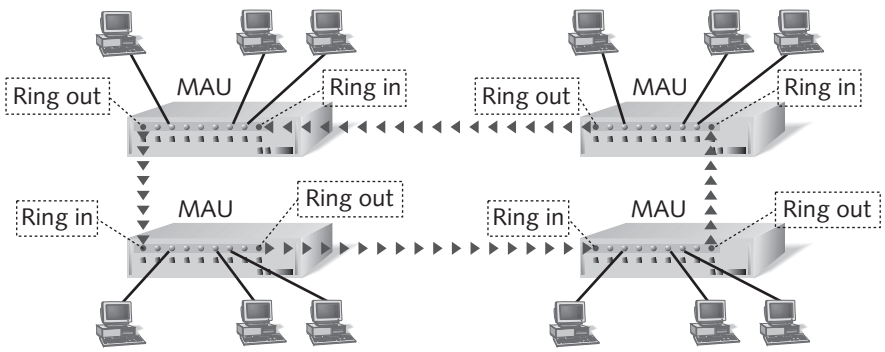

Figure 7-10 Tokens pass from one host to another in a unidirectional logical ring topology

Ethernet

Ethernet is by far the most widely used media access method today because it offers a nice balance between cost, speed, and ease of installation. Since Ethernet utilizes shared media between nodes, there are rules for sending packets of data to avoid collisions and protect the data. Though Ethernet collisions are unavoidable, you want to minimize their occurrence as much as possible to optimize available bandwidth, not waste it with excessive collisions. A large number of collisions can occur because there are too many users

on the network contending for bandwidth. Segmenting (subnetting) the network into separate, smaller networks joined together with a switch or router is one way of reducing traffic on an overcrowded network. The Institute for Electrical and Electronic Engineers (IEEE) has defined standards for Ethernet known as the **802.3 Standard**, which defines how to configure an Ethernet network as well as how elements in an Ethernet network interact with one another. By following the 802.3 standard, network equipment and network protocols can communicate properly.

Carrier Sense Multiple Access with Collision Detection (CSMA/CD) describes the method Ethernet devices use to negotiate access to the wire and retransmit in case of collisions. The sending host monitors the voltage level of the wire, and if no transmission is occurring, the host sends data. If two or more hosts determine that the network is clear and begin sending data at the same time, Collision Detection (CD) handles timely attempts to retransmit data for each host.

Although electrical signals on Ethernet travel at speeds nearing the speed of light, it still takes a finite amount of time for the signal to travel from one end of a large Ethernet network to another. In larger network designs, the signal quality begins to depreciate as segments exceed their maximum length. Ethernet hubs repeat the signal, extending the maximum length, and connect two or more Ethernet segments of any media type. A **hub** provides a universal link for devices in a network and sends all incoming data out to all ports (hence, to each node). There are several types of hubs, including:

- A **passive hub**, which provides a channel for the data, enabling it to go from one device (or segment) to another.
- An **intelligent hub** (or **managed hub**), which includes additional components that enable administrators to monitor the traffic passing through the hub and to configure each port in the hub.
- A **switching hub**, which reads the destination address of each packet and then forwards the packet to the correct port.

If the hub is attached to a backbone, then all computers can communicate with all the hosts on the backbone. A very important fact to note about hubs is that they only allow users to share Ethernet bandwidth. A network of hubs/repeaters is called a "shared Ethernet," meaning that all member hubs contend for data transmission on a single network (collision domain). The number of hubs in any one collision domain is restricted by the 802.3 rules. This means that individual members of a shared network will only get a percentage of the available network bandwidth.

Ethernet is governed by the "5-4-3 rule" of repeater placement. This rule means that the network can only have five segments connected; it can only use four repeaters; and of the five segments, only three can have users attached to them—the other two must be inter-repeater links. If the design of the network breaks these rules, then the timing guidelines will not be met and the sending node will resend that packet, resulting in lost packets and excessive resent packets. This can adversely affect network performance by slowing down the network and creating problems for applications.

Fast Ethernet

Fast Ethernet (IEEE 802.3u) offers higher transmission speeds than 802.3 Ethernet. (Fast Ethernet is also known as 100BaseT.) Fast Ethernet allows for fewer repeaters because the data travels so quickly that host NICs cannot always compensate for collision detection and consequent retransmissions in a timely manner. In Fast Ethernet networks, there are two classes of repeaters. Class I repeaters have a latency of 0.7 microseconds or less and are limited to one repeater per network. Class II repeaters have a latency of 0.46 microseconds or less and are limited to two repeaters per network. Fast Ethernet can be deployed to desktops and servers by installing Fast Ethernet NICs and using Fast Ethernet switches and repeaters. This standard raises the Ethernet speed limit from 10 Mbps to 100 Mbps with no changes to the existing cable structure or connectors. Most of today's networks have a mixture of standard Ethernet networks (10 Mbps) and Fast Ethernet (100 Mbps).

Full-duplex Ethernet

There is another variation of Ethernet called **full-duplex Ethernet**. By simply adding another pair of wires (total of six wires) and removing collision detection, you can double the connection speed. Hosts can simultaneously send and receive data similar to a telephone conversation in which both parties can speak at once. (Half-duplex would be more like a CB radio conversation.) In terms of Fast Ethernet, 200 Mbps of throughput is the theoretical maximum for a full-duplex Fast Ethernet connection. This type of connection is limited to a node-to-node connection and often links two Ethernet switches. Full duplex is just another method used to increase bandwidth to dedicated workstations or servers by doubling the bandwidth on a link, providing 20 Mbps for Ethernet and 200 Mbps for Fast Ethernet. To use full duplex, special NICs are installed in the computers and a switch is programmed to support full-duplex operation.

You can't use full duplex with a hub, because the nature of full duplex requires a dedicated connection. Therefore, you would have to use a switch (which provides full bandwidth to each port) instead of a hub (which shares bandwidth with other ports).

Gigabit Ethernet

Gigabit Ethernet is a newer version of Ethernet that supports data-transfer rates of 1 Gigabit (1000 megabits) per second. The first Gigabit Ethernet standard (802.3z) was ratified by the IEEE 802.3 Committee in 1998 and is defined by the frame format, the use of CSMA/CD, the use of full duplex, the use of flow control, and the management objects defined by the committee. Gigabit Ethernet is basically Ethernet, only faster. Most organizations use Gigabit Ethernet as a backbone technology and for server connections. Gigabit Ethernet is a future technology that promises a migration path beyond Fast Ethernet; the next generation of networks will support even higher data-transfer speeds. The first installations will require fiber optic media for long connections between buildings, and short copper links for connections between servers and hubs. Over time, as the market for workgroup and desktop Gigabit Ethernet services develops, customers will demand Gigabit links that are compliant with the installed base of Category 5 UTP wiring that is used for standard Ethernet. (Wiring specifics appear later in this chapter.)

Network Equipment

Repeaters allow networks to broaden the distance limitations of the cabling; however, repeaters support only a limited number of stations. As we saw earlier, with shared Ethernet, the likelihood of collision is higher as nodes are added to the shared collision domain. Using a bridge or switch to segment the traffic is a good way to resolve this problem. A switch can replace a hub or repeater and improve network performance. Bridges and switches allow LANs to expand considerably because they can maintain full Ethernet segments on each port. Bridges and switches can also filter network traffic so that traffic destined for the same network does not pass through the switch to another network. Traffic destined for another network is forwarded to the other network.

Bridges

Though there are several types of bridges, the basic function of a **bridge** is to connect separate networks. Bridges connect similar or different network types, like Ethernet and Token Ring. In this scenario, bridges map the Ethernet addresses of the nodes residing on each network segment and let only certain traffic pass through. When the bridge receives a packet, it determines both the destination host and the source host location. The bridge drops (filters) the packet if the source and destination are on the same segment. If the source and destination are on different segments, the bridge passes the packet on to the correct segment. This is why bridges are called "store-and-forward" devices; they analyze the Ethernet packet before making the decision to filter or forward. By using bridges to filter packets and regenerate forwarded packets, one can split a network into separate collision domains, allowing for more repeaters to be used in the total network design and achieving greater distances. In Figure 7-11, a message from a host on Segment 2 passes through the bridge to Segment 1 only when the message is

addressed to a host on Segment 1. This differs from a single segment connected to a repeater, which would repeat the message to all hosts.

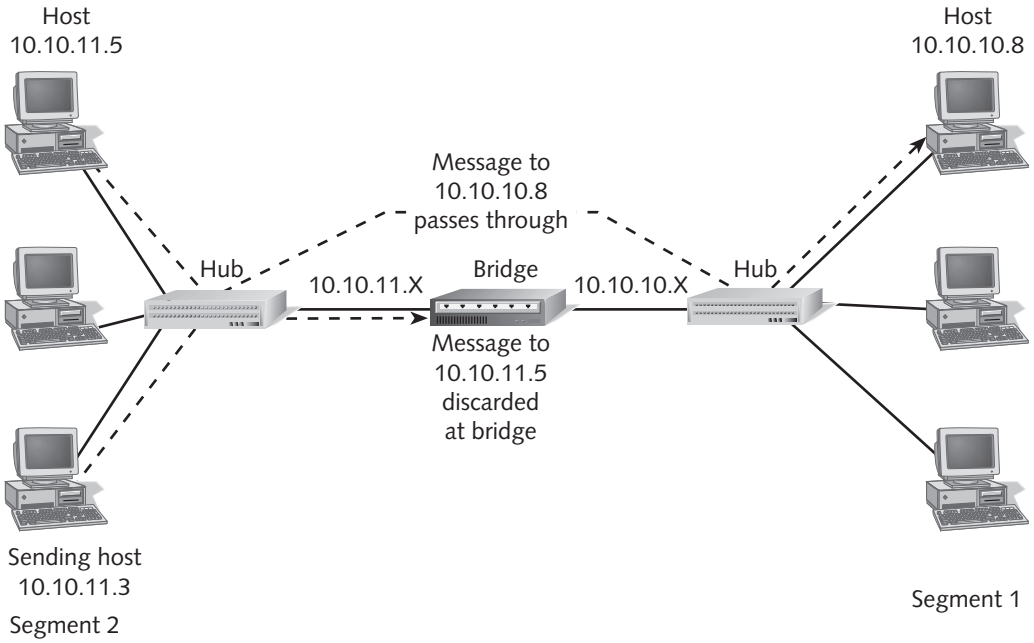

Figure 7-11 Only traffic destined for 10.10.10.8 passes through the bridge

Switches

Like bridges, there are several types of Ethernet switches, the details of which are beyond the scope of this book. Ethernet switches are an expansion of the concept in Ethernet bridging. **Switches** separate a network into collision domains so that network rules can be extended. Each of the segments attached to an Ethernet switch has a full 10 or 100 Mbps of bandwidth shared by fewer users, resulting in better performance (see Figure 7-12).

In addition to determining when to forward or filter a packet, Ethernet switches totally regenerate the packet, allowing each port on a switch to be treated as a complete Ethernet segment able to support the full length of cable as well as all of the repeater restrictions. Ethernet switches also recognize bad packets and instantly drop them from the network. This action keeps problems restricted to a single segment and prevents them from disrupting the rest of the network. Newer switches offer even higher bandwidth through Fast Ethernet, fiber, or ATM. These technologies link switches together or give added bandwidth to high-traffic servers.

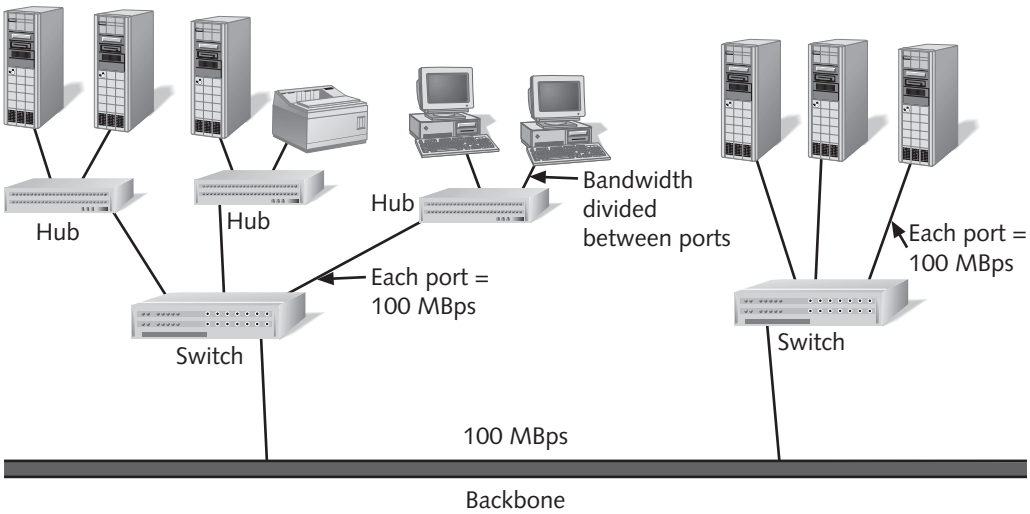

Figure 7-12 Switches offer full bandwidth to each port and forward or filter traffic similarly to a bridge

There is also another device that interconnects WAN links or the Internet to your network. This device is a CSU/DSU (Channel Service Unit/Digital (or Data) Service Unit). This device is necessary to connect the network to an external digital line such as a T-1 Internet line. (Sometimes the acronym is reversed and appears as DSU/CSU, but the meaning is the same.)

Routers

A **router** connects multiple networks using routing tables and routable protocols. Routers use headers and a forwarding table to determine where packets go, and communicate with other routers to calculate the best route between any two hosts. Routers determine whether to forward or filter a packet based on the IP address and subnet mask, which identifies the network to which a host belongs. The router filters a message destined for a host on the same network, and forwards messages destined for a host on a different network. While this functionality sounds similar to a bridge or switch, it differs in that a bridge or switch is not designed to enable networking over a large geographical area (such as a WAN), and routers are typically more capable of handling extremely high throughput. Also, because of the route calculations, routers can intelligently determine the best path from source to host over multiple routers. Switches and bridges do not include this kind of intelligence.

For example, in Figure 7-13, a user in Workgroup C sends a print job to a network printer located in Workgroup A. The message goes through Hub C to Router C. Router C analyzes the packet and determines that the message should go to Router A based on the message destination. Router C determines the best path to Router A, which might be to send the message directly to Router A instead of first passing it to Router B.

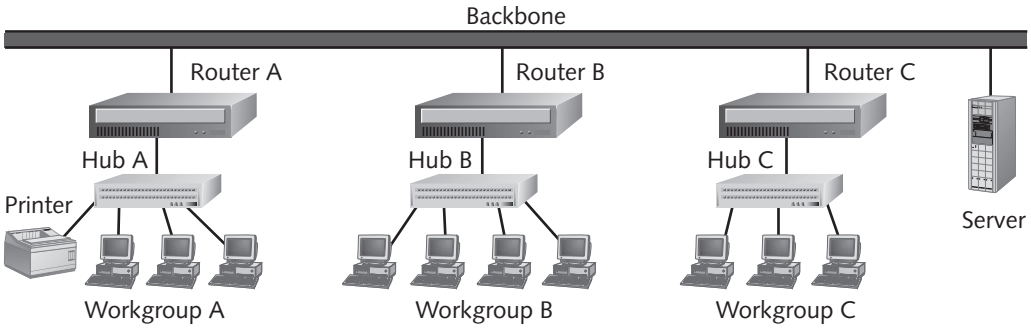

Figure 7-13 Routers determine the best route for sending messages

Network Cabling

An important part of designing and deploying a network is selecting the appropriate cable medium. The cable provides the conduit for all network communications. The network media access method standard of the NIC and the physical topology affect the cable choice and implementation. The most common types of LAN cabling in use today are coaxial (thinnet and thicknet), shielded twisted pair, unshielded twisted pair, and fiber optic.

Although this chapter primarily addresses physical cable media, there are wireless options that comply with the new Ethernet 802.11b standard to allow transmission at rates up to 11 Mbps using radio frequency within a specific radius.

Thicknet

Thicknet is based on the 10Base5 standard, which transmits data at 10 Mbps over a maximum distance of 500 meters (1640.4 feet). Thicknet is about 1 cm thick and has been used for network backbones because of its durability and maximum length. Many current thicknet implementations are being replaced by fiber optic media (addressed later in this chapter).

Thinnet

Thinnet is based on the 10Base2 standard (10 Mbps/Baseband transmission) that utilizes RG-58 A/U or RG-58 C/U 50 ohm coaxial cable with maximum segment lengths of 185 meters (606.9 feet). The RG-58 cable is less expensive and easier to install than the thicknet cable used for the 10Base5 standard, mostly because it is thinner (approximately 0.5 cm or .2 inch) and more flexible. While thinnet is gradually fading into networking history, you will still find it in several existing (but not new) network implementations. Cables in both the 10Base5 and 10Base2 systems interconnect with BNCs (British Naval Connectors) (see Figure 7-14). The network interface card in a

computer requires a T-connector to attach two cables to a NIC. A BNC barrel connector connects two cables. Any unused connection must have a 50 ohm terminator to prevent signal bounce.

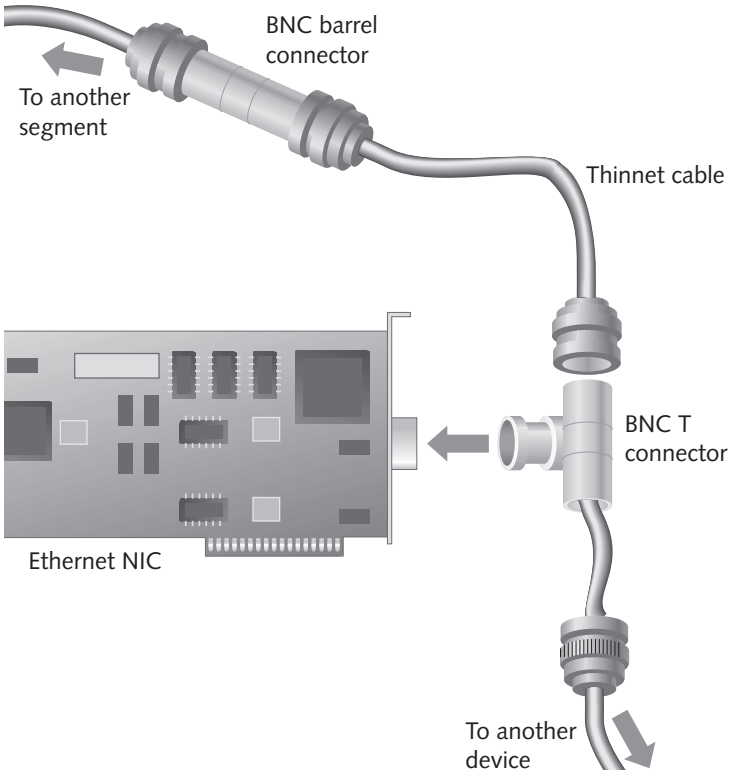

Figure 7-14 Thinnet uses several types of BNCs

Though RG-58 thinnet looks like TV cable, it is not, and the two cable types are not interchangeable because of differences in ohm impedance. Both thicknet and thinnet are 50 ohms, and TV cable is 75 ohms.

Shielded Twisted Pair (STP)

Shielded twisted-pair (STP) cable includes screened twisted-pair cable and foil twisted-pair cable and provides reliable connectivity. STP involves two copper wires, each encased in its own color-coded insulation, and then twisted together to form a "twisted pair." Multiple twisted-pairs are then packaged in an outer sheath to form the twisted pair cable (see Figure 7-15). The cable minimizes the possibility of **crosstalk** (intruding signals from an adjacent twisted pair or cable) by increasing the number of twists per inch. Early telephone signals were actually sent over a form of twisted-pair

cable, and almost every building today still uses twisted-pair cable to carry telephone and other signals. Although coaxial and fiber optic cable were developed to handle higher-bandwidth applications and support emerging technologies, twisted-pair cable has evolved so that it can now carry high-data-rate signals.

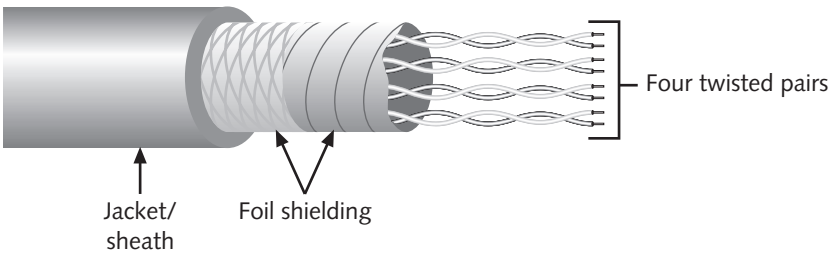

Figure 7-15 STP cable

STP cable encases the wires in a conducting metal shield to reduce the potential for EMI (as discussed in Chapter 2). Recall that electric motors, power lines, fluorescent lighting, and a variety of other devices cause EMI and, as a result, disruptions in network communications. STP cable effectively prevents radiation and blocks interference as long as the entire end-to-end link is shielded and properly grounded. The maximum length of STP is 100 meters (328.08 feet) with a speed of up to 500 Mbps. Although length becomes a problem with STP, it is inexpensive and easy to install.

STP cable can use several types of connectors, but only RJ-45 connectors are used in current networking contexts. **RJ-45 (registered jack-45)** is an eight-wire connector that connects Ethernet network devices. RJ-45 connectors look similar to the RJ-11 connectors that are used for connecting telephone equipment, but they are wider because they connect eight wires instead of four (see Figure 7-16).

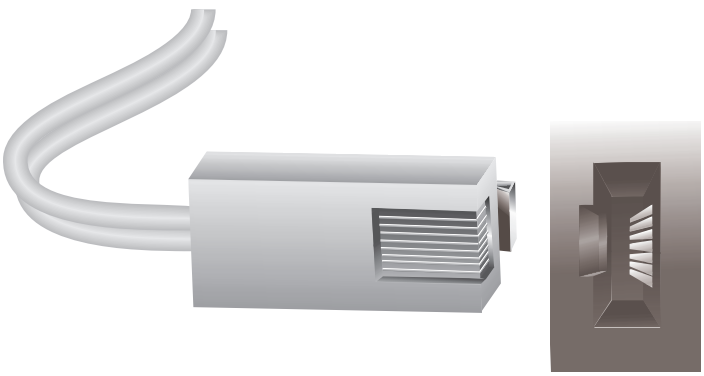

Figure 7-16 RJ-45 male and female connectors

Unshielded Twisted Pair (UTP)

Unshielded twisted-pair (UTP) cable, on the other hand, does not rely on physical shielding to block interference, but uses balancing and filtering techniques to reduce signal interference. Noise is induced equally on two conductors, which cancel out at the receiver. With correctly designed and manufactured UTP cable, this technique is easier to maintain than the shielding permanence and grounding of an STP cable. UTP cabling does not offer the high bandwidth or the protection from interference that coaxial or fiber optic cables do; however, millions of nodes are wired with UTP cable due to its lightweight, thin, and flexible nature. UTP cabling is a low-cost, manageable solution that is widely used for LANs and telephone connections. It is also quite adaptable and dependable, even for higher-data-rate applications. Like STP, UTP uses RJ-45 connectors. UTP has a maximum length of 100 meters (328.08 feet) and can support speeds of up to 100 Mbps.

Table 7-2 summarizes the various network cable types discussed so far.

Table 7-2 Network Cable Summary

Cable Type	Maximum Length	Maximum Speed
Thicknet (10Base5)	500 meters (1640.4 feet)	10 Mbps
Thinnet (10Base2)	185 meters (606.9 feet)	10 Mbps
Shielded twisted pair (STP)	100 meters (328.08 feet)	10, 100, 1000 Mbps depending on the category of cabling
Unshielded twisted pair (UTP)	100 meters (328.08 feet)	10, 100, 1000 Mbps depending on the category of cabling

UTP cable has evolved over the years, and different varieties are available for different needs. Improvements such as variations in the twists, individual wire sheaths, or overall cable jackets have led to the development of several standards for STP/UTP cabling. Here are some of the most popular specifications, each using four wire pairs:

- *Category 3 (Cat 3)*—Provides signal throughput up to 10 Mbps. Cat 3 is permissible for 10 Mbps Ethernet; however, the slightly more expensive Cat 5 is more common and reliable. You can also find Cat 3 cable in older 4 Mbps Token Ring network implementations.
- *Category 4 (Cat 4)*—Provides signal throughput up to 16 Mbps. Cat 4 is common in 16 Mbps Token Ring network implementations. Cat 5 is preferred over Cat 4 in most cases.
- *Category 5 (Cat 5)*—Provides signal throughput up to 100 Mbps. Cat 5 is the most common network cable in use today, and is used for Ethernet as well as other fast networking technologies. Cat 5 can be used for 10BaseT, 100BaseT, 1000BaseT, and Token Ring networking. Note that Enhanced Cat 5 (Cat 5e) is available for optimum reliability and performance. Cat 5e is only marginally more expensive yet reduces crosstalk by using a twisted string or plastic rib to

separates the pairs and keeps the wires in proper position along the whole length of the cable. Cat 5e also uses better insulation and a thicker jacket.

- *Category 6 (Cat 6)*—A heavily shielded implementation of twisted-pair cabling. Each of four pairs is wrapped in foil insulation, and even more foil wraps around the bundle of pairs. A fire-resistant sheath covers the outer foil layer. Cat 6 cabling provides excellent resistance to crosstalk and can support up to six times the throughput of Cat 5.

RJ-45 Pin Assignments

Cables normally connect to RJ-45 connectors in what is known as a "straight-through" configuration. This means that whatever color insulated wire is on the first pin on one end is the same color insulated wire on the first pin of the connector on the other end, and so on through all eight wires. Straight-through cabling is the most commonly used configuration and connects servers and workstations to wall connectors, hubs, and other network equipment.

While there are a few standards for determining the wiring order within the jack, it really doesn't matter as long as the wires on one end correspond exactly to the wires on the other end and the scheme is consistent throughout your organization. For reference sake, Figure 7-17 shows the Electronic Industries Association/Telecommunications Industry Association (EIA/TIA) 568A and 568B cable end standards. Generally speaking, 568B is more common in business and industry, and 568A is required in certain government installations. For straight-through cable, just make sure that both ends use the same configuration.

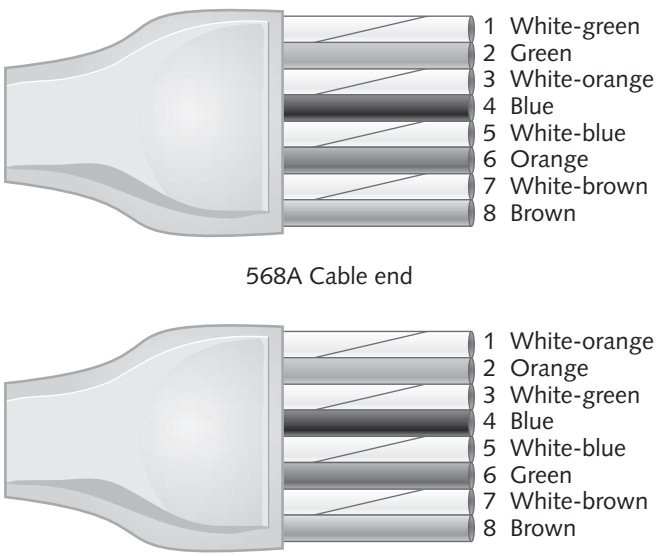

Figure 7-17 The EIA/TIA RJ-45 wiring standard

Crossover Cables

Crossover cables are useful in connecting two computers directly to one another. This is handy for quick connectivity without intermediate hubs or other network equipment as well as diagnostic contexts in which you want to isolate a host from the network and connect to it directly. Crossover cables are also required to daisy chain network devices. For example, crossover cables would interconnect a series of stackable hubs or switches (see Figure 7-18).

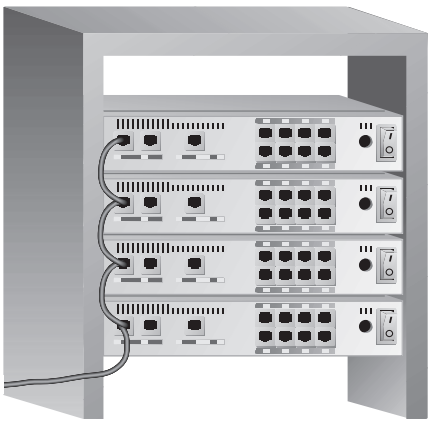

Figure 7-18 Crossover cables daisy chain network devices such as hubs

In a crossover configuration, you simply reverse the green and orange pairs. A simpler way to view this is to look at Figure 7-17 again. Recall that a straight-through cable uses the exact same pin configuration on both sides of the cable. For a crossover cable, simply use the 568A configuration on one end and the 568B configuration on the other.

Both a straight-through cable and crossover cable only use four of the eight pins: 1, 2, 3, and 6! You could goof the other wire assignments and not notice the difference.

How to Cut and Crimp RJ-45 Connectors

Although you can buy premade cables up to 25 feet (7.62 m) in length with guaranteed connectivity from manufacturers such as Belkin (*www.belkin.com*), you might find it worth your while to make your own cable under the following circumstances:

- You require lengths of cable longer than 25 feet (7.62 m).
- It is more economically feasible to purchase cable-making materials in bulk than premade cables.
- You want to create both standard and crossover cables.

- You want to choose your own color cable, useful for identifying cable runs to various parts of the building.

Many organizations choose orange cable to uniquely identify crossover cables.

To make your own cable, you will need the following materials:

- *Modular plug crimp tool*—These pliers allow you to crimp the connector onto the cable (see Figure 7-19) and often include sharp blades for stripping the jacket and insulation from cable. (Be careful of the blades!)

Figure 7-19 Most crimp tools include strippers and are able to crimp four to eight wires

- *Universal UTP stripping tool*—This optional tool is useful if you strip a lot of cable; it strips both round and flat cable (see Figure 7-20).

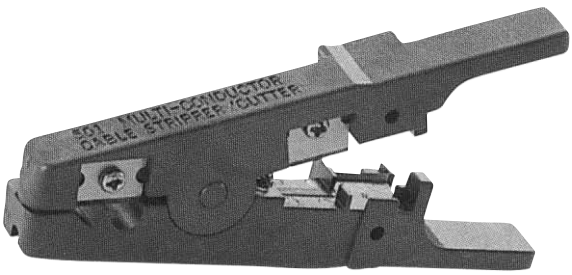

Figure 7-20 A universal UTP stripping tool

- *Diagonal cutters*—Use "diags" to cut cable off of the reel and to evenly trim pair ends during cable assembly (see Figure 7-21).

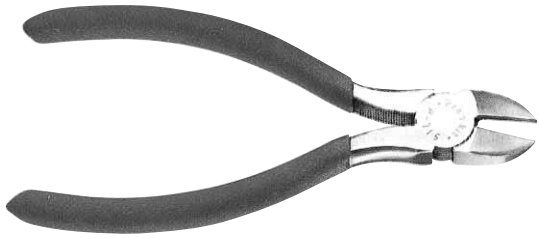

Figure 7-21 Diagonal cutters neatly cut and trim cable

- *Spooled cable*—UTP bulk cable is wound onto a spool inside a box (see Figure 7-22), and many manufacturers include an outlet from which to pull cable. There may also be a metering device so that you can measure the amount of cable pulled.

Figure 7-22 Spooled cable feeds through an outlet

- *Snagless boots*—Optionally, add a snagless boot (see Figure 7-23) to prevent the clip on RJ-45 cable from overextending and snapping off, which is a common problem when handling Cat 5 cable.

Figure 7-23 A snagless boot protects the RJ-45 clip

Cut and crimp your network cable as follows:

1. Pull the amount of cable you require. Cable is relatively inexpensive, so consider overestimating—too little cable requires you to start over, but seldom do you have too much.
2. Use the diags to cleanly cut the cable.
3. If you want to use snagless boots, thread two (one for each end) onto the cable. Make sure they are a few inches away from where you will crimp the ends to keep them out of the way.
4. Using the stripper or crimper, strip about an inch of the outer jacket off the end of the cable. Be careful not to cut into the insulation of the pairs inside. Although you can use the diags or some other tool to do the same task, I think you'll like the neatness of a stripper/crimper. Insert the end under the blade, lightly squeeze the crimper, and twist. This action cuts the outer jacket, which should then easily pull off.
5. Separate each of the four wire pairs and straighten them into the wiring order you choose (refer to Figure 7-17).
6. Evenly snip the end of the wires using the diags to expose a half inch of the wire pairs. (Originally, we cut one inch off of the jacket, but that was just to make handling the small wire pairs easier.) This should allow you to insert the wire pairs all the way into the end of the RJ-45 connector without leaving too much untwisted wire, which would cause crosstalk signaling problems. Also, half-inch exposure positions the jacketed portion of the cable under a plastic plug or wedge that clamps down onto the jacket.

Some sources recommend stripping the end of each individual wire, which is completely unnecessary and extremely time-consuming, as the crimper forces connection pins to pierce the insulation, making a solid connection.

7. Insert the wire ends all the way into the RJ-45 jack—be careful not to get the wires crossed or bent. Once you get the wires into their respective grooves in the jack, the rest is easy.
8. Confirm that the wiring order is correct.

9. Hold the cable (not the plug) and insert the plug into the crimper tool. If the crimper tool has plugs for multiple connectors, look for the one marked "8P."
10. Squeeze the crimper around the connector. Most crimpers are ratcheted so that the crimpers hold their place if you lose or need to readjust your grip partway through. The crimper includes two plungers that attach the connector (see Figure 7-24). The leftmost plunger depresses a plug or wedge onto the jacket to prevent strain from dislodging the connector. The second plunger depresses a two-pronged metal pin into each wire, creating the actual connection.

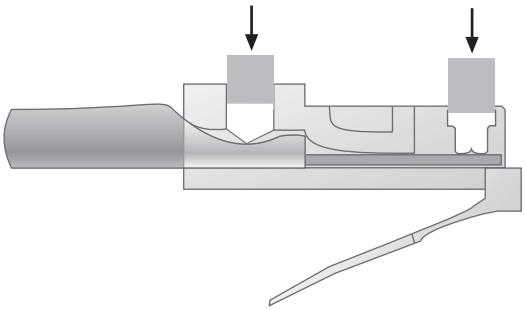

Figure 7-24 Two plungers attach the connector to the cable

11. Inspect the cable with the RJ-45 clip facing away from you. You should be able to see the wiring order through the clear plastic, as seen in Figure 7-25. If it's wrong, just use the diags to snip off the connector and start over again. (And don't feel bad—everybody goofs a plug sooner or later.) Also, give it a gentle tug; the connector shouldn't budge.

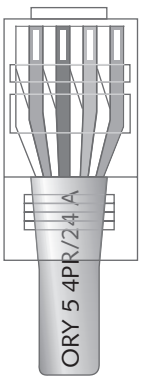

Figure 7-25 Inspect the wiring order

12. Repeat all steps for the other end of the cable.
13. Finally, if you included snagless boots, push them over the connectors.

Get more practice in making cables in Hands-on Project 7-1 at the end of the chapter.

10BaseT

The 10BaseT standard (also called twisted-pair Ethernet) uses a twisted-pair cable (such as Cat 5) with maximum lengths of 100 meters (328.08 feet). The cable is thinner and more flexible than the coaxial cable used for the 10Base5 or 10Base2 standards. If you were cabling for a 10BaseT network, the use of two pairs of Cat 3 wires would be sufficient; however, Cat 5 is a better choice in new implementations for purposes of signal quality and future upgrades.

100BaseT

The committee responsible for 100BaseT kept the 100 Mbps Ethernet standard as close to the original Ethernet definition as possible. Therefore, 100BaseT utilizes the Carrier Sense Multiple Access with Collision Detection (CSMA/CD) shared-media access method supported in earlier versions of Ethernet. The simplicity of this media access method might make it attractive to companies using traditional Ethernet. Because it is 10 times faster than Ethernet, it is often referred to as Fast Ethernet. Officially, the 100BaseT standard is called IEEE 802.3u. There are several different cabling schemes that can be used with 100BaseT, including the following:

- *100BaseT4* uses an extra two wires (four pairs) of normal-quality twisted-pair wires for use with Cat 3 UTP cable.
- *100BaseTX* Fast Ethernet uses two pairs of high-quality twisted-pair wires for use with Cat 5 UTP cable. The 100BaseTX standard has become the most popular due to its close compatibility with the 10BaseT Ethernet standard.
- *100BaseFX* is used with fiber optic cable that, for the most part, connects hubs and switches either between wiring closets or between buildings. 100BaseFX uses multimode fiber optic cable to transport Fast Ethernet traffic.

Fiber Optic Cabling

Fiber optic cable employs a technology that uses glass (or plastic) threads (fibers) to transmit data using light pulses from one end of the cable to the other. A fiber optic cable consists of a core of glass threads, each of which is capable of transmitting light pulses. A cladding surrounds the fibers, mirroring light back into the core. Plastic and braided Kevlar form the protective jacket (see Figure 7-26). The receiving end of the message converts the light signal to binary values. The maximum length is 25 km (15.5 miles) with speeds up to 2 Gbps. Fiber optics is not a particularly popular technology for workstations because it is so expensive. However, it is often used for a network backbone. Fiber cabling has been deployed as the primary media for campus and building backbones, offering high-speed connections between diverse LAN segments. Today, with progressively more complicated applications such as high-speed ISPs and e-commerce, optical fiber may soon become the primary media providing data to the desktop. Fiber has the largest bandwidth of any media available. It can transmit signals over the longest distance at the lowest cost, with the fewest repeaters and the least amount of repairs.

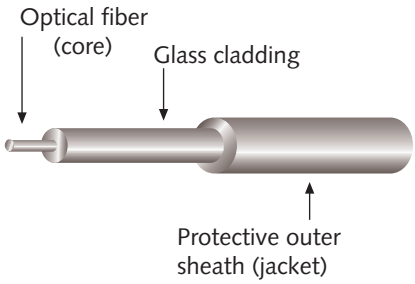

Figure 7-26 Fiber optic cable

Following are advantages of fiber optics:

- Fiber optic cables have a much greater bandwidth than metal cables to carry more data.
- Fiber optic cables are less susceptible to interference than metal cables.
- Fiber optic cables are much thinner and lighter than metal wires.
- Data can be transmitted digitally rather than in analog format.
- Fiber is immune to EMI and radio frequency interference (RFI). Because it does not conduct electricity, fiber optic cable can also be useful in areas where large amounts of EMI are present, such as on a factory floor.
- There are no crosstalk issues.
- Fiber is impervious to lightning strikes and does not conduct electricity or support ground loops.
- Fiber-based network segments can be extended 20 times farther than copper segments.
- Fiber cable cannot be tapped, so it's very secure.
- Fiber transmission systems are highly reliable.

Despite the numerous advantages of fiber, the cables are expensive to install, are more fragile than wire, and are difficult to split.

Network Adapters

Many NIC adapters comply with Plug and Play specifications so that they can be automatically configured without user intervention. On a non-Plug and Play system, the configuration must be done manually through a software program and/or DIP switches or jumpers, especially on older NICs. Network adapter cards are available to maintain practically every networking specification, including Fast Ethernet. Fast Ethernet NICs are often 10/100 Mbps capable and will automatically adapt to the proper speed. For

network operating systems, adapter-teaming techniques can be used to offer additional bandwidth and improved performance. **Adapter teaming** is the process of installing two or more network adapters in a server and then logically grouping them so that they appear to the operating system as a single network interface. There are several different types of teaming techniques, including Adaptive Fault Tolerance (AFT), Adaptive Load Balancing (ALB), and link aggregation.

Increase network performance to and from the server and the client by installing only 32-bit PCI network cards instead of 8- or 16-bit ISA network cards. The PCI cards can provide several times the network throughput because of the wider data path and more efficient use of the faster PCI bus.

Adaptive Fault Tolerance (AFT)

Adaptive Fault Tolerance (AFT) can provide an easy and effective method for increasing the availability of network server connections. By simply installing two or more server network adapters and configuring AFT, you have an emergency backup connection between the server and the network. If there is any problem with a cable, NIC, switch port, or hub port on the primary adapter, the secondary adapter can kick in within seconds to provide transparent recovery for applications and users. AFT can be configured with just two server adapters. Certain vendors and manufacturers provide automatic AFT support when multiple server adapters are configured.

A critical point of vulnerability on your network is corporate or departmental servers, where a failure or bottleneck can be disastrous to productivity. AFT technology usually supports up to four adapter teams with two to four adapters on each team. Depending on the vendor, AFT solutions can be made up of various adapter types and speeds as long as there is at least one primary server adapter in the team. The primary server adapter will generally pass its MAC and Layer 3 address to the failover adapter(s). This type of AFT requires the NIC vendor to provide a driver model that can take advantage of this technology.

Adaptive Load Balancing (ALB)

Adaptive Load Balancing (ALB) is a technique of guaranteeing a consistent level of high server throughput and transparent backup links by implementing multiple NICs and balancing the data transmission load across them. ALB is also known as asymmetric port aggregation. With Intel components, for example, you can use as many as four server adapters, connect them to a switch, and configure them to work as a team for an aggregate throughput of up to 400 Mbps with Fast Ethernet adapters or 8 Gbps with Gigabit Ethernet adapters. All of the adapters in a team must be connected to a switch, and the team is assigned a single network address. (ALB is designed to work with a switch, whereas AFT works with either a hub or a switch.) With ALB, all of the traffic moving from the network server is automatically balanced between up to four links. This can

assure fast throughput with no need to restructure or reconfigure the network. An illustration of ALB can be seen in Figure 7-27.

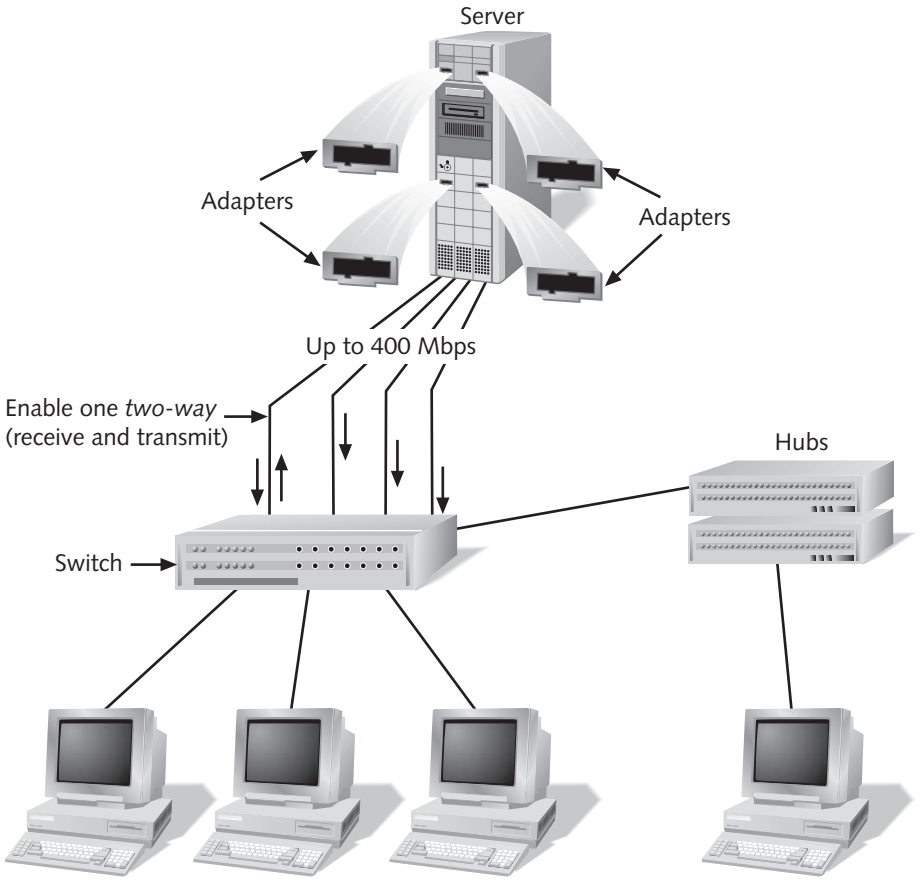

Figure 7-27 Adaptive Load Balancing increases throughput to and from the server

ALB offers a simpler and better way to move more data faster through the server by enabling each adapter to essentially add another 100 Mbps link, or channel. ALB also delivers the same fault-tolerance benefits as AFT technology, because if one link fails, the others will continue to provide network connectivity. Therefore, with ALB technology, it is no longer necessary to segment the network if the server link becomes a bottleneck. Instead, you can eliminate the bottleneck quickly and easily by installing two adapters in your server and configuring ALB with the driver software. This requires no client configuration, and clients do not have to be routed to communicate with each other. In addition, traffic is balanced along all of the server adapters. As with AFT, an intelligent, adaptive agent is included with the software driver. This driver dynamically manages the server adapter team and evenly distributes the load among them by constantly analyzing the

traffic flow from the server. One channel within an ALB team carries traffic to and from the server while the others carry traffic from the server only. This load balancing of server traffic assures that all users enjoy the same network response from the server. By taking advantage of ALB, a four-link configuration can yield an aggregate throughput of approximately 400 Mbps.

An intelligent adaptive agent in the driver will also continuously analyze the traffic flow from the server and distribute the packets based on destination addresses. Load balancing can occur only on routed protocols (IP and NCP IPX). Multicast/broadcast and nonrouted protocols such as NetBEUI and Microsoft IPX are transmitted only over the primary adapter.

For AFT and ALB technologies, you can usually choose a primary and secondary role for selected adapters. The primary adapter will carry the majority of the traffic. With AFT, it will be the only adapter used until that link fails. With ALB and nonroutable protocols, it will be the only adapter used. It will also be the only adapter used for broadcast and multicast traffic.

Link Aggregation

Link aggregation is the process of combining multiple adapters into a single channel to provide bandwidth greater than the base speed of the adapter (10, 100, or 1000 Mbps). Link aggregation works only across multiple source address/destination address pairs. ALB mode provides aggregation in both directions with aggregation-capable switches. Intel's version of link aggregation, for example, must be used with an Intel Express switch capable of link aggregation and must match the capability of the switches. All adapters share a single MAC address. Link aggregation can be implemented in Windows 2000 and Windows NT 4.0, NetWare 4.1 and above, SCO UnixWare (10/100), and Linux.

Networking with a Modem Pool

Although networking is mostly LAN/WAN communication between hosts, remote access (the ability to connect to a particular server or the LAN over a phone line or Internet connection) is increasingly a concern for networked organizations as well. Most users have Internet access at home, and there are methods with which users can securely connect directly to the LAN through their ISP. However, there will likely be a need for users to directly dial into the LAN using conventional voice telephone lines, or POTS (Plain Old Telephone Service, also known as PSTN, Public Switched Telephone Network). Some network policies require dialing in directly because dedicated phone lines have less chance of exposure to mischievous Internet users. However, dialing in requires a modem on both ends. On the user end, this is an inexpensive device that's possibly already integrated into the motherboard of the computer. On the server end, you must provide a receiving modem to receive dozens or hundreds of callers at once (even thousands of callers at once if you run an ISP).

Servicing so many users quickly presents a problem. Because you typically install a modem in a free expansion slot on the computer, you probably only have room for one to three modems on a server. If you service dozens or hundreds of dial-up users, this type of modem installation is not feasible. Instead, you can install a **modem pool**, which uses one or more external physical devices that represent several modems (see Figure 7-28).

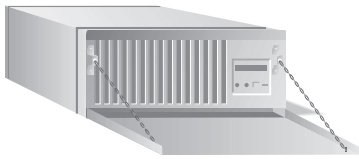

Figure 7-28 An analog modem pool

Starting from the user's remote dial-up attempt, Figure 7-29 illustrates the various equipment to which the user transparently connects as well as the function of each component. Here's how a modem pool works:

1. The user dials in, usually using a standard analog modem.
2. The dial-in attempt passes through normal Telco channels.
3. The dial-in attempt rings into the network modem pool, usually over a T-1 connection (a high-bandwidth phone connection capable of voice and data transmission).
4. A modem in the modem pool accepts the call. Note that although each modem in the pool operates independently, multiple users dial in using the same telephone number. Modem pool intelligence logically separates each call.
5. The call is useless until it reaches the LAN or server. Depending on the vendor, the modem pool device connects to an adapter card (which performs the serial processing on behalf of the server), or connects directly to a port on a network device such as hub, switch, or router.
6. The server either restricts the call to itself (somewhat unusual but sometimes implemented as a security measure to protect the LAN) or, more commonly, connects the call to the LAN. Once connected to the LAN, the user experiences the same connectivity as when locally connected to the LAN (except at slow modem speed).

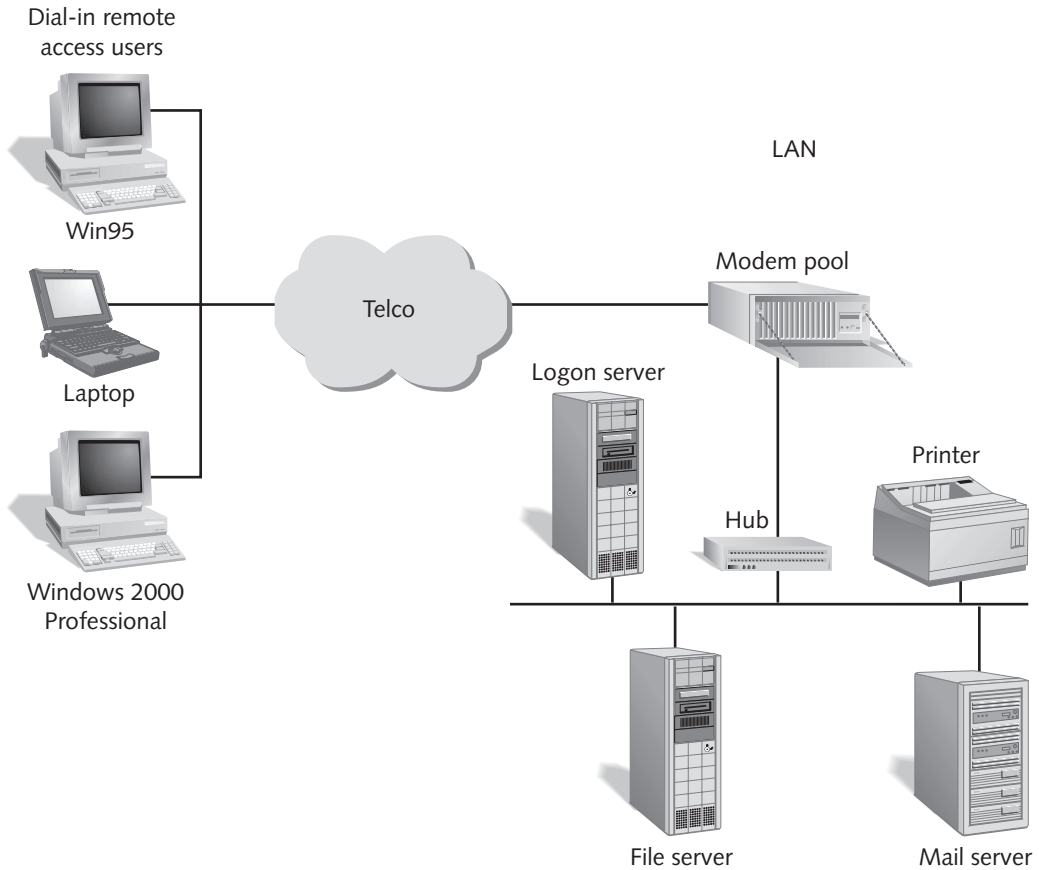

Figure 7-29 Connecting remote users to the LAN

The modem pool can be a desktop device, but larger implementations usually implement one or more rack-mounted devices. A modem pool is either analog or digital.

- *An analog modem pool* has a very basic chassis construction that provides expansion slots for multiple modems (usually 16). The chassis, in turn, connects to an adapter card in the server. This implementation is fading away because a digital implementation offers much higher density and more sophisticated configuration.
- *A digital modem pool* is a device with circuitry and intelligence to comprise the modems. Some vendors' modem pools (for example, Equinox/Avocent) accept SIMM modules (not to be confused with memory modules) that represent multiple modems. For example, a digital modem pool with four installed SIMM modules, each representing six modems, has a total of 24 modems. Digital modem pools can accept standard analog modem calls or digital ISDN calls. You can daisy chain multiple devices attached to the adapter card to increase density.

Besides accepting dial-in modem and ISDN calls, many modem pools also dial out. This is a common implementation to increase security to verify that only connections to a specific phone number are accepted. A fax pool uses a very similar construction.

Many modem pool implementations (such as the Cisco AS5800 Universal Access Server) are capable of modem pool partitioning. This is useful when you want to use separate phone numbers to connect various modem protocols such as V.90, V.34, and K56flex (see Figure 7-30).

Modems	Modems in pool	Assigned number
56K modems	24	555-1111
V.34 modems	24	555-2222
V.90 modems	24	555-3333

Figure 7-30 Separate phone numbers connect various modem protocols

Vendor software allows you to configure connections. For example, you can limit bandwidth, the number of simultaneous connections, the length of time a user can connect, and so forth.

Connecting the modem pool device is usually very simple. You install the device in a standard 19-inch wide rack opening (see Chapter 4), connect the incoming communication line, and then connect the modem pool device to the network equipment (see Figure 7-31).

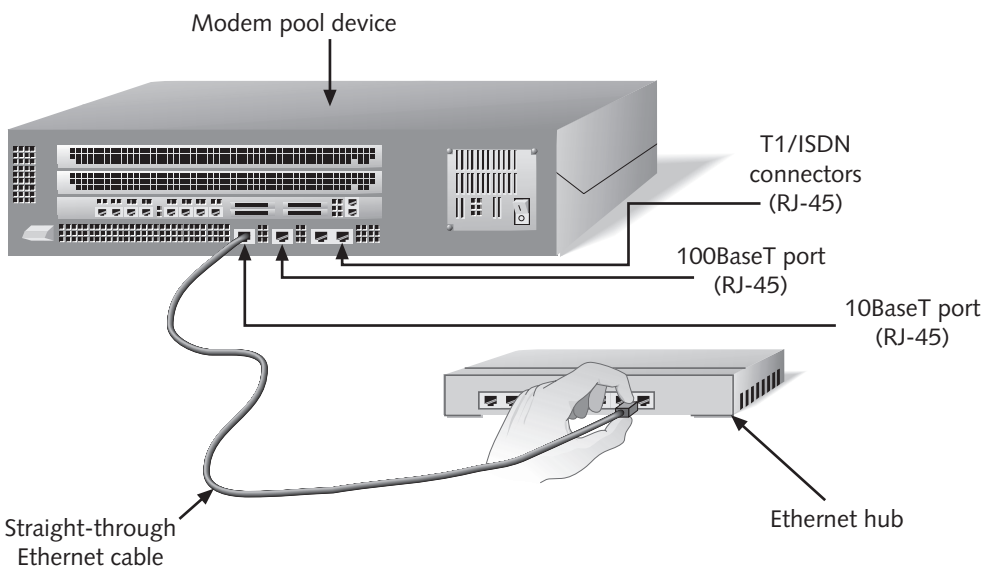

Figure 7-31 A common modem pool connection scenario

Chapter Summary

- Bus networks are comparatively inexpensive and easy to install for small networks. However, it is more difficult to troubleshoot and locate a break in the cable or a faulty machine on a bus.
- When a node on a Token Ring network needs to transmit data and receives an empty token, it holds on to the token and records its own address, the destination address, and the message into the token before passing it on to the next station. As long as the token is not marked as empty, all of the other stations will not transmit messages. The token passes from one node to the next until it completes a full circuit and reaches the originating station, where the message is discarded and the token is again marked as empty.
- Ring topology has the following advantages: It can span greater distances than other network types; it offers low signal deterioration; collisions are low because only the station that holds the token can transmit; and it is good for a small network requiring high transmission speeds or for larger networks where each station has a comparable workload.
- Ring topology has the following disadvantages: It can be tricky to trace a problem on the cable segment if the LAN is large; each station's attached network interface must be continually active and the failure of a single station will halt a unidirectional ring network; it is complex to configure and requires relatively expensive hardware for each computer to interface with the network; and transmission delays tend to be long, even with moderate traffic levels.
- In the star topology, all nodes receive the same signal, reducing effective bandwidth, and the network device can become a bottleneck because all data must pass through it. Standard twisted-pair Ethernet networks using 10Base5, 10BaseT, or 100BaseTX technology commonly use the star topology.
- Star topology has the following advantages: A single failed node does not adversely affect the rest of the network; it is relatively straightforward to install and manage; isolating and repairing bad segments is easier; and it offers good capacity for network growth.
- Star topology has the following disadvantages: It needs a lot more cabling than bus or ring networks; and the entire network becomes ineffectual if the central network device fails.
- Ultimately, network protocols are only a means to transport the network message to the node's physical address, known as the MAC (Media Access Control) address, which globally and uniquely identifies a network device.
- NetBEUI is not a routable protocol. It is not efficient in larger networks because it frequently rebroadcasts to locate other nodes on the network since it does not cache previously located nodes and does not use name resolution services such as DNS or WINS.

- IPX/SPX includes a caching mechanism so that it is not necessary to rebroadcast to locate recently accessed nodes, but it does not utilize name resolution services such as DNS or WINS.
- Administrators consider TCP/IP a single protocol, though it is actually a suite of protocols. TCP/IP is more difficult than other protocols to plan and configure, although it is also scaleable and routable, which is why it is the protocol of the Internet and most enterprise networks.
- TCP/IP uses a unique IP address to identify each node. An IP address appears as four sets of digits, each separated by a dot. An IP address must be globally unique if the host is on the Internet. Each IP address also requires a subnet mask that, when compared against the IP address, identifies the specific network to which the host belongs.
- Administrators use a subnet mask to divide a range of IP addresses into multiple smaller networks because they might not need all the available IP addresses on a single network. Instead, you might split the IP address range among several separate networks. Administrators also subnet their networks to split up a collision domain.
- Improve network throughput to and from the server using multihoming—multiple network connections to a single server. You can also use port aggregation to combine multiple server network connections into a single, logical network connection with bandwidth that is multiplied times the number of ports.
- Ping tests connectivity at different stages between the host and destination to determine the point of failure at which a packet is dropped, and also tests basic networking connectivity.
- IPCONFIG is a Microsoft utility that displays a wide variety of IP configuration data for a Windows 98/ME/NT/2000 system, including the IP address, subnet mask, and default gateway.
- ARP displays the resolution between the IP address and physical (MAC) address on the network interface card by building a table as IP addresses resolve to MAC addresses.
- TRACERT is a trace routing utility that works like Ping but shows the actual router hops taken to reach the remote host.
- NETSTAT with the –r option shows routing tables.
- Token Ring networks, which operate at 4 or 16 Mbps, use the ring topology.
- Ethernet is the most widely used access method because it offers a nice balance between cost, speed, and ease of installation. Since Ethernet utilizes a shared medium between nodes, there are rules for sending packets of data to avoid conflicts and protect the data.
- With Carrier Sense Multiple Access with Collision Detection (CSMA/CD), the sending host monitors the voltage level of the wire, and if no transmission is occurring, the host sends data. If two or more hosts determine that the network is clear and begin sending data at the same time, Collision Detection (CD) handles timely attempts at retransmitting data for each host.

- A hub provides a universal link for devices in a network and sends all incoming data out to all ports and, hence, to each node.
- A network of hubs/repeaters is called a "shared Ethernet," meaning that all members of the network are contending for data transmission on a single network (collision domain).
- The Ethernet network can only have five segments connected; it can only use four repeaters; and of the five segments, only three can have users attached to them—the other two must be inter-repeater links. This is called the "5-4-3 rule."
- Fast Ethernet (100BaseT) offers transmission speeds up to 100 Mbps over Cat 5 cable and RJ-45 connectors.
- By simply adding another pair of wires (total of six wires) and removing collision detection, you can double the connection speed, creating full-duplex Ethernet, which must use switches.
- Gigabit Ethernet is a full-duplex media access method supporting transfer rates up to 1 Gbps.
- A bridge connects similar or different network types, such as Ethernet and Token Ring. Bridges drop (filter) the packet if the source and destination are on the same segment. If the source and destination host are on different segments, bridges pass the packet on to the correct segment.
- Switches separate a network into collision domains so that network rules can be extended. Each of the segments attached to an Ethernet switch has a full 10 or 100 Mbps of bandwidth shared by fewer users, resulting in better performance.
- A router connects multiple networks using routing tables and routable protocols. Routers use headers and a forwarding table to determine where packets go, and communicate with other routers to calculate the best route between any two hosts.
- Thinnet is based on the 10Base2 standard (10 Mbps/Baseband transmission) that utilizes RG-58 A/U or RG-58 C/U 50 ohm coaxial cable with maximum segment lengths of 185 meters (606.9 feet).
- Shielded twisted-pair (STP) cable involves two copper wires, each encased in its own color-coded insulation, which are twisted together to form a "twisted pair." Multiple twisted pairs are then packaged in an outer sheath to form the cable. The cable minimizes the possibility of crosstalk by increasing the number of twists per inch. STP uses RJ-45 connectors.
- Unshielded twisted-pair (UTP) cable uses RJ-45 connectors. UTP has a maximum length of 100 meters (328.08 feet) and can support speeds up to 100Mps. UTP uses balancing and filtering techniques through media filters to reduce signal interference. Noise is induced equally on two conductors, which cancel out at the receiver.
- Category 3 (Cat 3) cable provides signal throughput up to 10 Mbps and is common in older 4 Mbps Token Ring network implementations. Category 4 (Cat 4) cable

provides signal throughput up to 16 Mbps and is common in 16 Mbps Token Ring network implementations. Category 5 (Cat 5) cable provides signal throughput up to 100 Mbps and is used for Ethernet as well as other fast networking technologies. Cat 5 can be used for 10BaseT, 100BaseT, 1000BaseT, and Token Ring networking.

- Cat 5 cables normally connect to RJ-45 connectors in what is known as a "straight through" configuration, in which the color insulated wire on the first pin on one end is the same color insulated wire on the first pin of the connector on the other end, and so on through all eight wires.
- Crossover cables connect two computers directly to one another and also daisy chain network devices. A crossover cable uses the 568A configuration on one end and the 568B configuration on the other.
- You probably want to make your own cable if you require lengths of cable longer than 25 feet (7.62 m), if it is more economically feasible to purchase cable-making materials in bulk than premade cables, if you want to create both standard and crossover cables, or if you want to choose your own color cable (useful for identifying cable runs to various parts of the building).
- The 10BaseT standard (also called "twisted-pair Ethernet") uses a twisted-pair cable (such as Cat 5) with maximum lengths of 100 meters (328.08 feet).
- 100BaseT utilizes the CSMA/CD shared-media access method supported in earlier versions of Ethernet. The simplicity of this media access method might make it attractive to companies using traditional Ethernet. Because it is 10 times faster than Ethernet, it is often referred to as Fast Ethernet.
- 100BaseT networking uses 100BaseT4, 100BaseTX (the most common implementation), or 100BaseFX cabling schemes.
- Fiber optic cable uses glass (or plastic) threads (fibers) to transmit data using light pulses from one end of the cable to the other, and consists of a core of glass threads, each of which is capable of transmitting light pulses. A cladding surrounds the fibers, mirroring light back into the core. Plastic and braided Kevlar form the protective jacket. The maximum length is 25 km (15.5 miles) with speeds up to 2 Gbps.
- The following are characteristics of fiber optic cable: much greater bandwidth than metal cable to carry more data; less susceptible to interference than metal cable; much thinner and lighter than metal cable; data can be transmitted digitally rather than in analog format; it is immune to EMI and RFI; it can be useful in areas where large amounts of EMI are present; there is no crosstalk; it is impervious to lightning strikes and does not conduct electricity or support ground loops; network segments can be extended 20 times farther than copper segments; and it cannot be tapped, so it's very secure.
- A modem pool can be a desktop device, but larger implementations usually have one or more rack-mounted devices. A modem pool falls under one of two categories: analog (a chassis for multiple physical modem adapters) or digital (a device with circuitry and intelligence to comprise the modems).

KEY TERMS

802.3 Standard — An Institute for Electrical and Electronic Engineers (IEEE) networking standard that defines the rules for configuring an Ethernet network as well as determining how elements in an Ethernet network interact with one another. By following 802.3, network equipment and network protocols can communicate properly.

adapter teaming — Installing two or more network adapters in a server and then logically grouping them so that they appear to the operating system as a single network interface.

Adaptive Fault Tolerance (AFT) — Installing two or more server network adapters to provide an emergency backup connection between the server and the network. If there is any problem with a cable, NIC, switch port, or hub port on the primary adapter, the secondary adapter can kick in within seconds to provide transparent recovery for applications and users.

Adaptive Load Balancing (ALB) — A technique of guaranteeing a consistent level of high server throughput and transparent backup links by implementing multiple NICs and balancing the data transmission load across them. ALB is also known as "asymmetric port aggregation."

ARP (Address Resolution Protocol) — Displays the resolution between the IP address and physical (MAC) address on the NIC by building a table as IP addresses resolve to MAC addresses. You can also modify the ARP cache and table entries.

backbone — A larger, common avenue through which data transfers take place from smaller lines connected to it.

bridge — A network device that connects separate networks together. Bridges connect similar or different network types, such as Ethernet and Token Ring.

bus topology — A network topology in which nodes link together in series where each node is connected to a common backbone cable.

collision—An event that results when two nodes transmit packets to the network at the same time.

collision domain — A network boundary in which multiple nodes could potentially attempt to access the network at the same time.

crosstalk — Intruding signals from an adjacent twisted pair or cable.

data rate — The actual quantity of data transferred within the limitations of the bandwidth.

Fast Ethernet — A 100 Mbps Ethernet implementation, also known as 100BaseT.

fiber optic cable — Technology that uses glass (or plastic) threads (fibers) to transmit data using light pulses. The receiving end of the message converts the light signal to binary values. The maximum length is 25 km (15.5 miles) with speeds up to 2 Gbps.

full-duplex Ethernet — The addition of another pair of wires (total of six wires) to Ethernet cable and removing collision detection to double the connection speed. Hosts can simultaneously send and receive data similar to a telephone conversation in which both parties can speak at once. (Half-duplex would be more like a CB radio conversation.)

Gigabit Ethernet — Supports data transfer rates of 1 Gigabit (1000 megabits) per second.

hub — A networking device that provides a universal link for devices in a network and sends all incoming data out to all ports (hence, to each node).

intelligent hub (or **managed hub**)—A hub that allows administrators to monitor the traffic passing through the hub and configure each port in the hub.

IPCONFIG — A Microsoft utility that displays a wide variety of IP configuration data for a Windows 98/ME/NT/2000 system, including the IP address, subnet mask, and default gateway and other information.

IPX/SPX (Internetwork Packet Exchange/Sequence Packet Exchange) — The default Novell protocol implementation for all versions of NetWare until 5.x, which can also use TCP/IP.

link aggregation — Combining multiple adapters into a single channel to provide bandwidth greater than the base speed of the adapter (10, 100, or 1000 Mbps). Link aggregation works only across multiple source address/destination address pairs.

MAC (Media Access Control) address — A globally unique identifier found on each NIC.

media access method — A method to place the data packets transmitted from the NOS software to the physical network device (such as a NIC) and then to the wire.

modem pool — One or more external physical devices that represent several modems.

multihomed — Computers using NICs with multiple ports or multiple NICs to increase effective network throughput.

multistation access unit (MAU) — A networking device that looks much like a hub except that it includes an RI (ring in) and RO (ring out) port. Tokens still pass from one host to the next in a logical ring.

NetBEUI (NetBIOS Enhanced User Interface) — A small, fast protocol optimized for small networks.

NETSTAT — A command-line networking utility that shows TCP/IP protocol statistics using any of several options. One of the most useful options is –r, which shows the routing table. This is useful in verifying the efficiency of the routing tables.

node — An active device connected to a network, such as a computer or printer, or networking equipment such as a hub, switch, or router.

passive hub — A standard hub that simply receives signals and repeats them out to all ports.

Ping (packet internet groper) — An all-purpose utility for verifying that a remote host is accessible by sending small packets of data to which an accessible host responds.

port aggregation — Similar to multihomed computers, port aggregation uses software to combine multiple ports from the server into a single logical connection to the network but with bandwidth that is multiplied times the number of ports.

ring topology — A network topology in which all of the nodes are connected in a closed, single, logical communication loop.

round-trip time (RTT) — The time it takes for a Ping packet to reach its destination and return to the source.

RJ-45 (registered jack-45) — An eight-wire connector that connects Ethernet network devices.

router — A network device that connects multiple networks using routing tables and routable protocols. Routers use headers and a forwarding table to determine where packets go, and communicate with other routers to calculate the best route between any two hosts. Routers determine whether to forward or filter a packet based on the IP address and subnet mask, which identifies the network to which a host belongs. The router filters a message destined for a host on the same network, and forwards messages destined for a host on a different network.

shielded twisted pair (STP) — Network cable that consists of two copper wires, each encased in its own color-coded insulation, which are twisted together to form a "twisted pair." Multiple twisted pairs are then packaged in an outer sheath to form the twisted-pair cable.

star topology — A network configuration in which all of the nodes connect to a central network device such as a hub or switch. All nodes receive the same signal, reducing effective bandwidth, and the central network device can become a bottleneck because all data must pass through it.

switch — A networking device that separates a network into collision domains so that network rules can be extended. Each of the segments attached to an Ethernet switch has a full 10 or 100 Mbps of bandwidth shared by fewer users, resulting in better performance.

switching hub — A hub that reads the destination address of each packet and then forwards the packet to the correct port.

subnet mask — A series of network identification numbers which, when compared against the IP address, identifies the specific network to which the host belongs.

TCP/IP (Transmission Control Protocol/Internet Protocol) — A suite of protocols in common use on most networks and the Internet.

thinnet — Based on the 10Base2 standard (10 Mbps/Baseband transmission), networking cable that utilizes RG-58 A/U or RG-58 C/U 50 ohm coaxial cable with maximum segment lengths of 185 meters (606.9 feet).

thicknet — Based on the 10Base5 standard, which transmits data at 10 Mbps over a maximum distance of 500 meters (1640.4 feet). Thicknet is about 1 cm thick and has been used for backbone media because of its durability and maximum length.

topology — The geometric configuration of devices, nodes, and cable links on a network. Topologies define how nodes connect to one another.

TRACERT — A command-line trace routing utility that works like Ping but shows the actual router hops taken to reach the remote host.

transceiver — A device in a Token Ring node that repeats the network signal to move it around the ring.

token passing — A method of collision avoidance that prevents two nodes from transmitting messages at the same time.

unshielded twisted pair (UTP) — Network cabling that does not rely on physical shielding to block interference (as does STP), but uses balancing and filtering techniques to reduce signal interference. Noise is induced equally on two conductors, which cancel out at the receiver.

Review Questions

1. Which network topology does a network backbone use?
   a. star
   b. bus
   c. hub
   d. ring
2. When can a node in the ring topology transmit a message?
   a. when the node receives an empty token
   b. when the node receives a token with the parity bit set to 1
   c. when the node receives a token with the parity bit set to 0
   d. any time it receives a token
3. What do you call a node on a TCP/IP network?
   a. Ping
   b. plug
   c. host
   d. gateway
4. What is the unit of measurement that corresponds to 1000 bits per second?
   a. Mbps
   b. Kbps
   c. Tbps
   d. Gbps
5. Which of the following network topologies best describes a network configuration where all of the nodes are connected to a centrally located hub or switch and is commonly used for standard twisted-pair Ethernet networks using 10BaseT or 100BaseTX?
   a. ring
   b. bus
   c. star
   d. mesh
6. Which of the following is characteristic of NetBEUI?
   a. routable
   b. caches location of other nodes
   c. no configuration
   d. complex configuration

7. IPX/SPX is the default protocol of which NOS?
   a. Windows NT 4.0
   b. NetWare versions before 5.0
   c. OS/2
   d. Linux
8. What is the purpose of a subnet mask?
   a. identifies the host's IP network
   b. uniquely identifies the host
   c. filters packets destined for the same network
   d. forwards packets destined for a different network
9. Which of the following are valid network diagnostic tools that are common to most operating systems? (Choose all that apply.)
   a. Ping
   b. TraceIP
   c. ARP
   d. NETSTAT
10. Why might you want to subnet the network?
    a. to break up the collision domain
    b. to give each department its own network
    c. to increase collisions
    d. to reduce traffic on the Token Ring network
11. Which one of the following is *not* a recognized version of the Ethernet standard?
    a. Fast Ethernet
    b. Gigabit Ethernet
    c. Token Ethernet
    d. Full-duplex Ethernet
    e. Switched Ethernet
12. Which one of the following cable types consists of four twisted pairs of copper wire terminated by RJ-45 connectors, supports frequencies up to 100 MHz and network speeds up to 100 Mbps, and can be used for 10BaseT, 100BaseT, and Token Ring networking?
    a. Category 3
    b. Category 4
    c. Category 5
    d. Category 2

13. Which one of the following is *not* a valid version of the 100BaseX specifications?
    a. 100BaseT4
    b. 100BaseTX
    c. 100BaseFX
    d. 100BaseF4
14. What is the purpose of IPCONFIG?
    a. to display name resolution results
    b. to display routing tables
    c. to display host IP configuration data
    d. to display router hops to a destination node
15. What is the difference between a MAU and an Ethernet hub?
    a. An Ethernet hub has an RI and RO port, and a MAU does not.
    b. An Ethernet hub is only found on a ring network, and a MAU only on a bus network.
    c. An Ethernet hub sends data in both directions in the ring, and a MAU sends data in only one direction.
    d. A MAU has an RI and RO port, and an Ethernet hub does not.
16. Which of the following is characteristic of a collision domain?
    a. Token Ring media access method
    b. Ethernet media access method
    c. A MAU arbitrates collisions.
    d. Any network with more than two nodes qualifies as a collision domain.
17. 10BaseT Ethernet networks use which of the following rules?
    a. 5 (segments) 4 (hubs) 3 (three of five segments can have users)
    b. 7 (segments) 6 (hubs) 5 (five of five segments can have users)
    c. 6 (segments) 5 (hubs) 4 (four of six segments can have users)
    d. use as many segments, hubs, and users as necessary within acceptable performance standards
18. A crossover cable:
    a. uses the same wiring order on both ends
    b. uses exactly the opposite wiring order on both ends
    c. connects a computer to a hub
    d. reverses orange and green pairs

19. How does a hub address a message destined for another network?
    a. It forwards the message directly to the other network.
    b. It repeats the signal indiscriminately through all ports.
    c. It directs the message only to the port of the specific destination host.
    d. It forwards the packet to other members of the ring.
20. What is the most common type of network cable in use today?
    a. thinnet
    b. Cat 3
    c. Cat 4
    d. Cat 5

HANDS-ON PROJECTS

Project 7-1

In this project, you will construct a straight-through cable and a crossover cable. You will need two or more RJ-45 jacks, about three feet of Cat 5 cable, diagonal cutters, and a crimper. You will use the cables you make in Project 7-2.

1. Measure about three feet (one meter) of cable from the spool.
2. Use the diags to cleanly cut the cable.
3. Using the crimper, strip about an inch of the outer jacket off the end of the cable. (If the crimper does not have blades for this purpose, you will need a stripper.) Be careful not to cut into the insulation of the pairs inside. Insert the end under the blade, lightly squeeze the crimper, and twist. This action cuts the outer jacket, which should then easily pull off.
4. Separate each of the four wire pairs and straighten them into the 568B wiring order (refer to Figure 7-17).

Evenly snip the end of the wires using the diags to expose a half inch of the wire pairs. This should allow you to insert the wire pairs all the way into the end of the RJ-45 connector.

5. Insert the wire ends all the way into the RJ-45 connector, being careful not to get the wires crossed or bent. Grooves in the connector help to guide the wires.
6. Confirm that the wiring order is correct.
7. Hold the cable (not the plug) and insert the plug into the crimper tool. If the crimper tool has plugs for multiple connectors, look for the one marked "8P."

8. Tightly squeeze the crimper around the connector.
9. Inspect the cable with the RJ-45 clip facing away from you. You should be able to see the wiring order through the clear plastic (as shown in Figure 7-25). If it's wrong, just use the diags to snip off the connector and start over again. When finished, give it a gentle tug—the connector shouldn't budge.
10. Repeat all steps for the other end of the cable.
11. Create a crossover cable using all the same steps, except that this time in Step 4 you will use the 568B connection on one side and the 568A connection on the other.

Project 7-2

In this project, you will construct a star topology network using a hub, the cables you made in Project 7-1, and at least two nodes.

1. Plug the power supply into one hub and make sure the power comes on (you should see one or more LED lights). Plug a straight-through network cable into one of the ports. Try to avoid using the last port, as some hubs link that port to the uplink so that you can only use either the last port or the uplink but not both.
2. Connect any Windows 2000 node to the hub with the other end of the straight-through cable.
3. Turn on the Windows 2000 computer.
4. Plug the power supply into a second hub and make sure the power comes on.
5. Plug a straight-through network cable into one of the ports, again avoiding the last port.
6. Connect another Windows 2000 node to the second hub with the other end of the straight-through cable. Coordinate with another student if using their computer.
7. Turn on the second Windows computer.
8. Using your crossover cable, connect the uplink ports between the two hubs.

Project 7-3

In this project, you will configure the IP address on the two Windows 2000 computers from Project 7-2 and verify their connectivity with Ping. This project requires you to log on as a local Administrator. Ask your instructor for the appropriate user name and password.

1. After the Windows computers have booted, verify that a link light appears on the network card, signifying connectivity.
2. Click **Start**, point to **Settings**, and click **Network and Dial-up Connections**. The Network and Dial-up Connections dialog box appears.
3. Right-click **Local Area Connection**, and click **Properties**.

4. In the Properties dialog box, click **Internet Protocol (TCP/IP)** and then click **Properties**. (Be careful not to uncheck Internet Protocol (TCP/IP).)
5. Configure the IP address to **10.10.10.X** where X is any number between 1 and 254 of your choosing. Configure the subnet address to 255.255.255.0.
6. Verify connectivity using the Ping utility. Ping the IP address of another connected node from a command prompt by typing **ping 10.10.10.X** where X is the address of the other node. You should see four replies if the nodes are properly configured and connected.
7. Close the Command Prompt window.

Project 7-4

In this project, you and a teammate will create a share and transfer a large quantity of files, and then observe the collisions on the hub(s). Perform the actions below on both computers.

1. On one of the Windows computers using Windows Explorer, create a folder from the C drive, **C:\Transfer**.
2. On the other computer, share the CD-ROM drive by right-clicking it and choosing **Sharing** from the context menu. Share the folder with default permissions.

In actual practice, sharing the CD-ROM on a server is not advisable because it opens security vulnerabilities.

3. Insert a CD-ROM with several files or the CD-ROM that came with this book.
4. On the first computer, use Windows Explorer to type **\\IPAddress** into the address bar at the top, where **IPAddress** is the IP address of the other computer. A list of shares appears in the details (right) pane of Windows Explorer.
5. Double-click the shared CD-ROM. You should see several files and/or folders. Select them all and copy them to your own C:\Transfer folder created in Step 1. Try to time the transfer so that both computers download contents of the CD-ROM from one another at the same time.
6. Observe the hub. The lights should blink constantly during the transfer, indicating collisions during the transfer. This is normal for an Ethernet network, even with only two hosts.
7. When the transfers are complete, delete the C:\Transfer folder and close Windows Explorer.

Project 7-5

In this project, you will reconfigure the computer's TCP/IP settings and connect to the Internet. Then, you will use the TRACERT command to trace the route to a few web sites.

1. Connect one of the Windows computers used in the other projects in this chapter to a line that has access to the Internet (your instructor will guide you).
2. Access the TCP/IP settings and configure them to receive an IP address automatically (from a DHCP server). If necessary, reboot the computer.
3. Browse to *www.ox.ac.uk*. This is the home page for Oxford University in London, England.
4. Open a command prompt. In the command prompt window, type **tracert *www.ox.ac.uk***.
5. Observe the number of hops required to access the site, which vary depending upon your physical location and Internet service provider. TRACERT counts up to 30 hops. The page you viewed in Step 3 came to you across all the routers you now observe in the TRACERT results.
6. Close the Command Prompt and browser windows.

Project 7-6

In this project, you will draw a network diagram that involves all three major topologies: bus, ring, and star.

1. On a separate sheet of paper, draw all three major network topologies, and interconnect them using any series of hubs, MAUs, bridges, and routers. There are scores of possible configurations that you can devise.
2. Share your network diagram with a classmate. At any point in the two diagrams, draw in a connecting device that might connect your two networks. For example, you could use a router or bridge to attach two bus topologies, or you could use a bridge to attach a Token Ring topology to an (Ethernet) star topology.

CASE PROJECTS

1. As a network consultant, you have been asked to analyze the network at Edna's Home Made Pies factory. The company has grown considerably since its inception six years ago, and users regularly complain about poor network performance, particularly in terms of logon response in the morning (all users report to work at 8:00 AM) and database performance. The owner, Edna, asks you to analyze the network and server placement within the network, and then to make recommendations.

   Upon arriving at the factory, you locate the server room, which has a wiring closet in the same room, and view all network equipment. Here are your findings:

   - A Novell NetWare server processes the user logons. It has a single 4 Mbps Token Ring network adapter.
   - A Windows NT 4.0 server runs the SQL database. Edna hires a database administrator to regularly optimize the database, and you cannot find any problems with the database itself. This server is on a 16 Mbps Token Ring network adapter on the same ring as the NetWare server.
   - A bridge appears to attach the server room servers on a Token Ring network to the user's Ethernet star topology. Last year, Edna had another consultant install the user workstations with Ethernet network adapters capable of up to 100 Mbps, but slow demand for pies restricted the budget to perform any more improvements.

   Given these findings, what would you recommend?

2. You are at a conference a few miles away from your workplace when your pager indicates a network emergency. Apparently, a network intern has done something to the physical network that has cut off several workstations from an Ethernet segment, and has taken a Linux server offline. When you arrive back at work, you inspect the physical connections and find the following:

   - Network cables daisy chain a series of stackable hubs. None of the cables are loose, and the power is on to each hub. Nevertheless, all hosts attached to hub C can only communicate with one another, and no hosts on any of the other hubs.
   - The Linux server seems to be operating properly, and the physical network connection to a switch is not loose. Nevertheless, you do not see a link light on the switch for this Linux server connection. Other connections to the switch show a green link light.

   Given these findings, what could be causing the problems?